


CONNEXIONS



The Interoperability Report

November 1993

Special Issue: INTEROP Europe 93 Companion

Volume 7, No. 11

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Global Internet Exchange.....	2
RIPE NCC.....	10
Profile: EUnet.....	15
MICE Multimedia.....	22
OSI applications on CLNS....	26
DFS.....	29
RENATER.....	32
APPC and APPN.....	34
Host Resources MIB.....	42
Announcements.....	44

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Mountain View, CA 94040-1219, USA. +1 415-941-3399. Fax: +1 415-949-1779. Toll-free: 1-800-INTEROP. E-mail: connexions@interop.com.

Copyright © 1993 by Interop Company. Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the *ConneXions* logo are registered trademarks of Interop Company.

ISSN 0894-5926

From the Editor

Welcome to INTEROP Europe 93 and to this special issue of *ConneXions—The Interoperability Report*. This edition contains articles directly and indirectly related to the conference, with pointers to sessions and tutorials at the end of each article.

The Internet is growing rapidly on a global scale. In order for such a large system to work well, a “sane” interconnection strategy must be in place.



25-29 October 1993 • CNIT, Paris La Défense, France

Our first article describes the *Global Internet Exchange* (GIX), a pilot effort to provide common routing exchange points, allowing pairs of networks to implement agreed upon routing policies. The article is by Bernhard Stockman who has been instrumental in the design and implementation of several European and international networking efforts.

When network operators troubleshoot routing and connectivity problems, they depend on a number of low-level tools which need a lot of non-local information to function properly. Since external routing depends on more than one operator, a means to exchange information about routing configurations and policies is necessary. This is the purpose of the RIPE NCC's *Routing Registry* database as outlined by Daniel Karrenberg.

Next we look at *EUnet*, one of Europe's most extensive Internet-related service organizations, connecting 27 countries with coverage from Iceland to Vladivostok and from the Arctic Circle to Northern Africa. The article is by Glenn Kowack, EUnet's Chief Executive.

Many efforts are underway to design multimedia communication systems. A variety of existing standards and technologies have been integrated in the MICE project, a joint effort between several academic and research institutions. MICE is described here by Jon Crowcroft of UCL.

Colin Robbins and Paul Barker describe their experiences with running OSI applications over CLNS networks. Such “tales from the trenches” are important as implementors attempt to deploy new technologies.

New technologies and standards also depend on some amount of “advocacy” to succeed. In this issue, you will find two articles which promote a certain approach: one on DFS and another on APPC/APPN. *ConneXions* does not endorse any particular technical solution, but rather we welcome alternative points of view, and regard articles as starting points for discussions.

France's new high-speed network, RENATER, is described in an article by Christian Michau, starting on page 32.

Our final article deals with the Host Resources MIB, a component of the SNMP network management framework. The article is by Steve Waldbusser of Carnegie Mellon University.

I would also like to draw your attention to another publication which is being distributed at INTEROP Europe 93. This month's issue of *TRIB-UNIX*, the publication of the Association of French UNIX Users, contains reprints of several articles that have previously appeared in *ConneXions*.

Global Connectivity: The Global Internet Exchange (GIX)

by Bernhard Stockman, SUNET

Introduction

The Internet has evolved from a single top level *Administrative Domain* (AD) using non hierarchical routing protocols such as EGP for inter-AD routing to an extensive collection of several transit ADs with different policies.

The routing protocols have been subsequently improved and the *Border Gateway Protocol* (BGP) can act as a basic toolbox for this type of environment. However, today's routing technology is based exclusively on destination address which very much limits flexibility. Another limiting factor is the rapid growth of the Internet in terms of the number of unique paths rather than networks. This will only increase as wide-scale deployment of BGP takes place.

At the IEPG/IETF meetings in November 1991, the issues of neutral interconnection points were initially discussed. At the March 1992 informal meeting, the idea was discussed in a day-long meeting in which goals and basic concepts were shared and at which a consensus began to emerge. At the June 1992 meeting in Tokyo, a detailed proposal was discussed and refined. After these meetings, and following discussions, a paper was produced proposing a GIX (*Global Internet eXchange*). The GIX would provide common routing exchange points, allowing pairs of networks to implement agreed upon routing policies and being independent of US Internet backbone structures.

Visions and Goals

The GIX concept is based on the vision of the Ubiquitous and Homogeneous Internet giving maximal connectivity and maximal flexibility. Maximal connectivity is where anyone can talk to anyone anywhere at any given time. Maximal flexibility is where every organization that participates can easily set their own rules and easily implement them.

The ambition is to make connections to this Internet from new communities of interest simple and straight forward. There should be no restriction on connections, but each network is autonomous in the sense that it is free to impose its own regional restrictions. In summary the goals are:

- To maximize global connectivity, i.e., to provide a way for new networks to attach simply and by that have paths to all other Internet connected networks.
- To implement a scalable management structure. The most important requirement is to provide stable, reliable, policy-based routing, with the policies being enforced by the networks at their attachment point rather than within the transit backbones.
- To allow networks to make their own decisions about transport, especially in cases where several different transit paths are possible.

Functionality

The functionality to be considered is:

- *Connectivity*: Connectivity to the Internet is global in the sense that once connected it should in principle be possible to reach all other Internet connected networks. There are several possible limitations on connectivity. Limits should only be imposed near users and not within transit backbones. Any limits imposed by usage constraints on intercontinental transit networks, are undesirable in the long run.

- *Transit*: Considering different demands on transit capacity a variety of options is needed to serve this diversity of transit. Transit capacity should be provided by the open market and by that be market driven with a decentralized management and a heterogeneous infrastructure.
- *Routing*: Routing is a global concern and depends on cooperation among Internet connected networks. For such collaboration to work there is a need for global guidelines and standards giving technical symmetry, stability, and manageability.

Requirements

The requirements of the GIX can be summarized as follows:

- *Scalability*: The size of the Internet, measured either in traffic or number of connected networks, is growing exponentially, and engineering and operations of the GIX must scale.
- *Manageability*: The worldwide Internet must be as well-managed as one under a single administration, and yet maintain the autonomy of each constituent network.
- *Accountability*: Any successful worldwide connectivity structure must be accountable.
- *Timeliness*: The cost of staying with the status quo, or the cost of delay, is very high. Maintaining the integrity of the Internet requires prompt action.

The GIX model

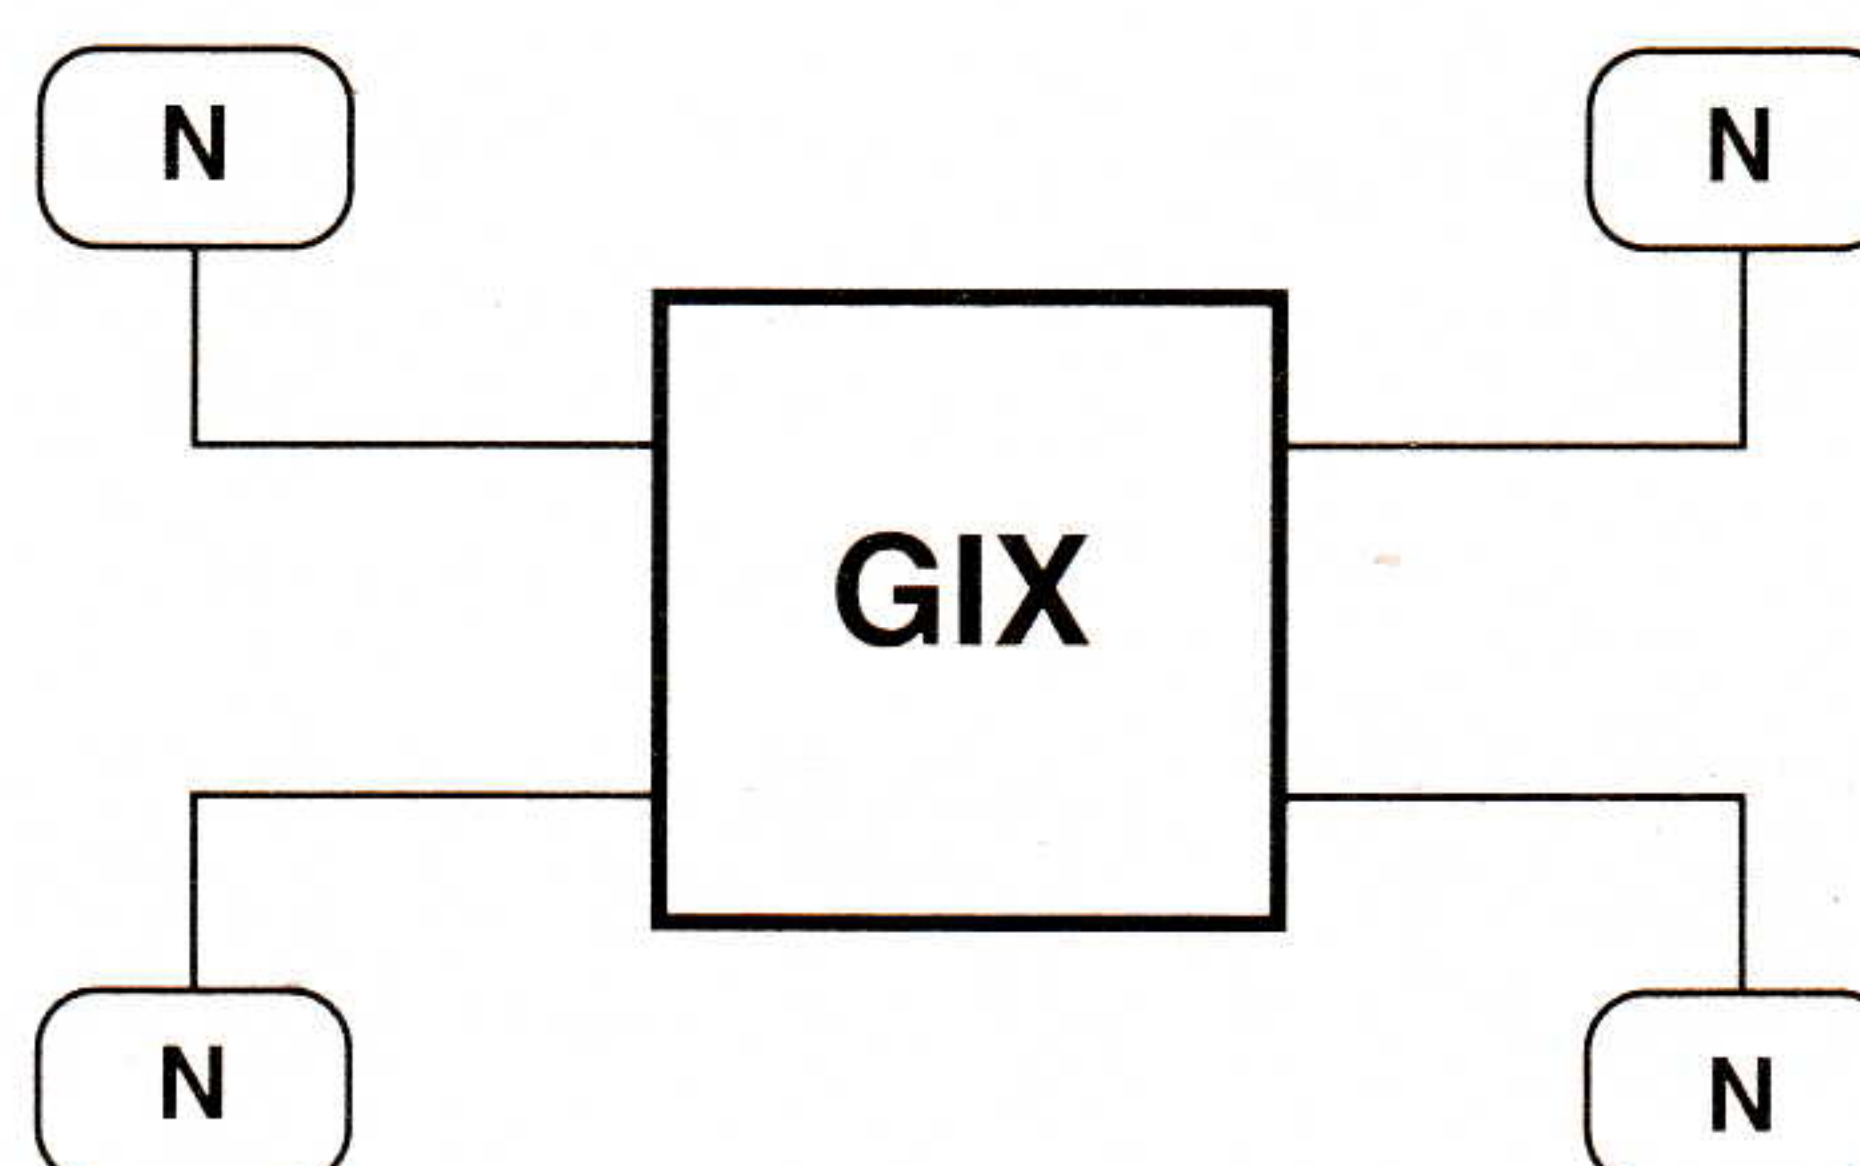
The GIX concept includes three basic functions:

- A physical layer-2 structure where networks can connect their routers and peer with any other routers connected.
- A *Routing Registry* serving as a neutral repository of registered routing policies.
- A *Route Server* which is a pseudo-router running on a host directly connected to the physical GIX.

Alternatives

Various models have been discussed. All with the requirements of providing capacities for routing, switching, transit and management. In the following diagrams the boxes marked GIX represent something giving routing, switching and transit capacity and the clouds marked N are community of interest networks.

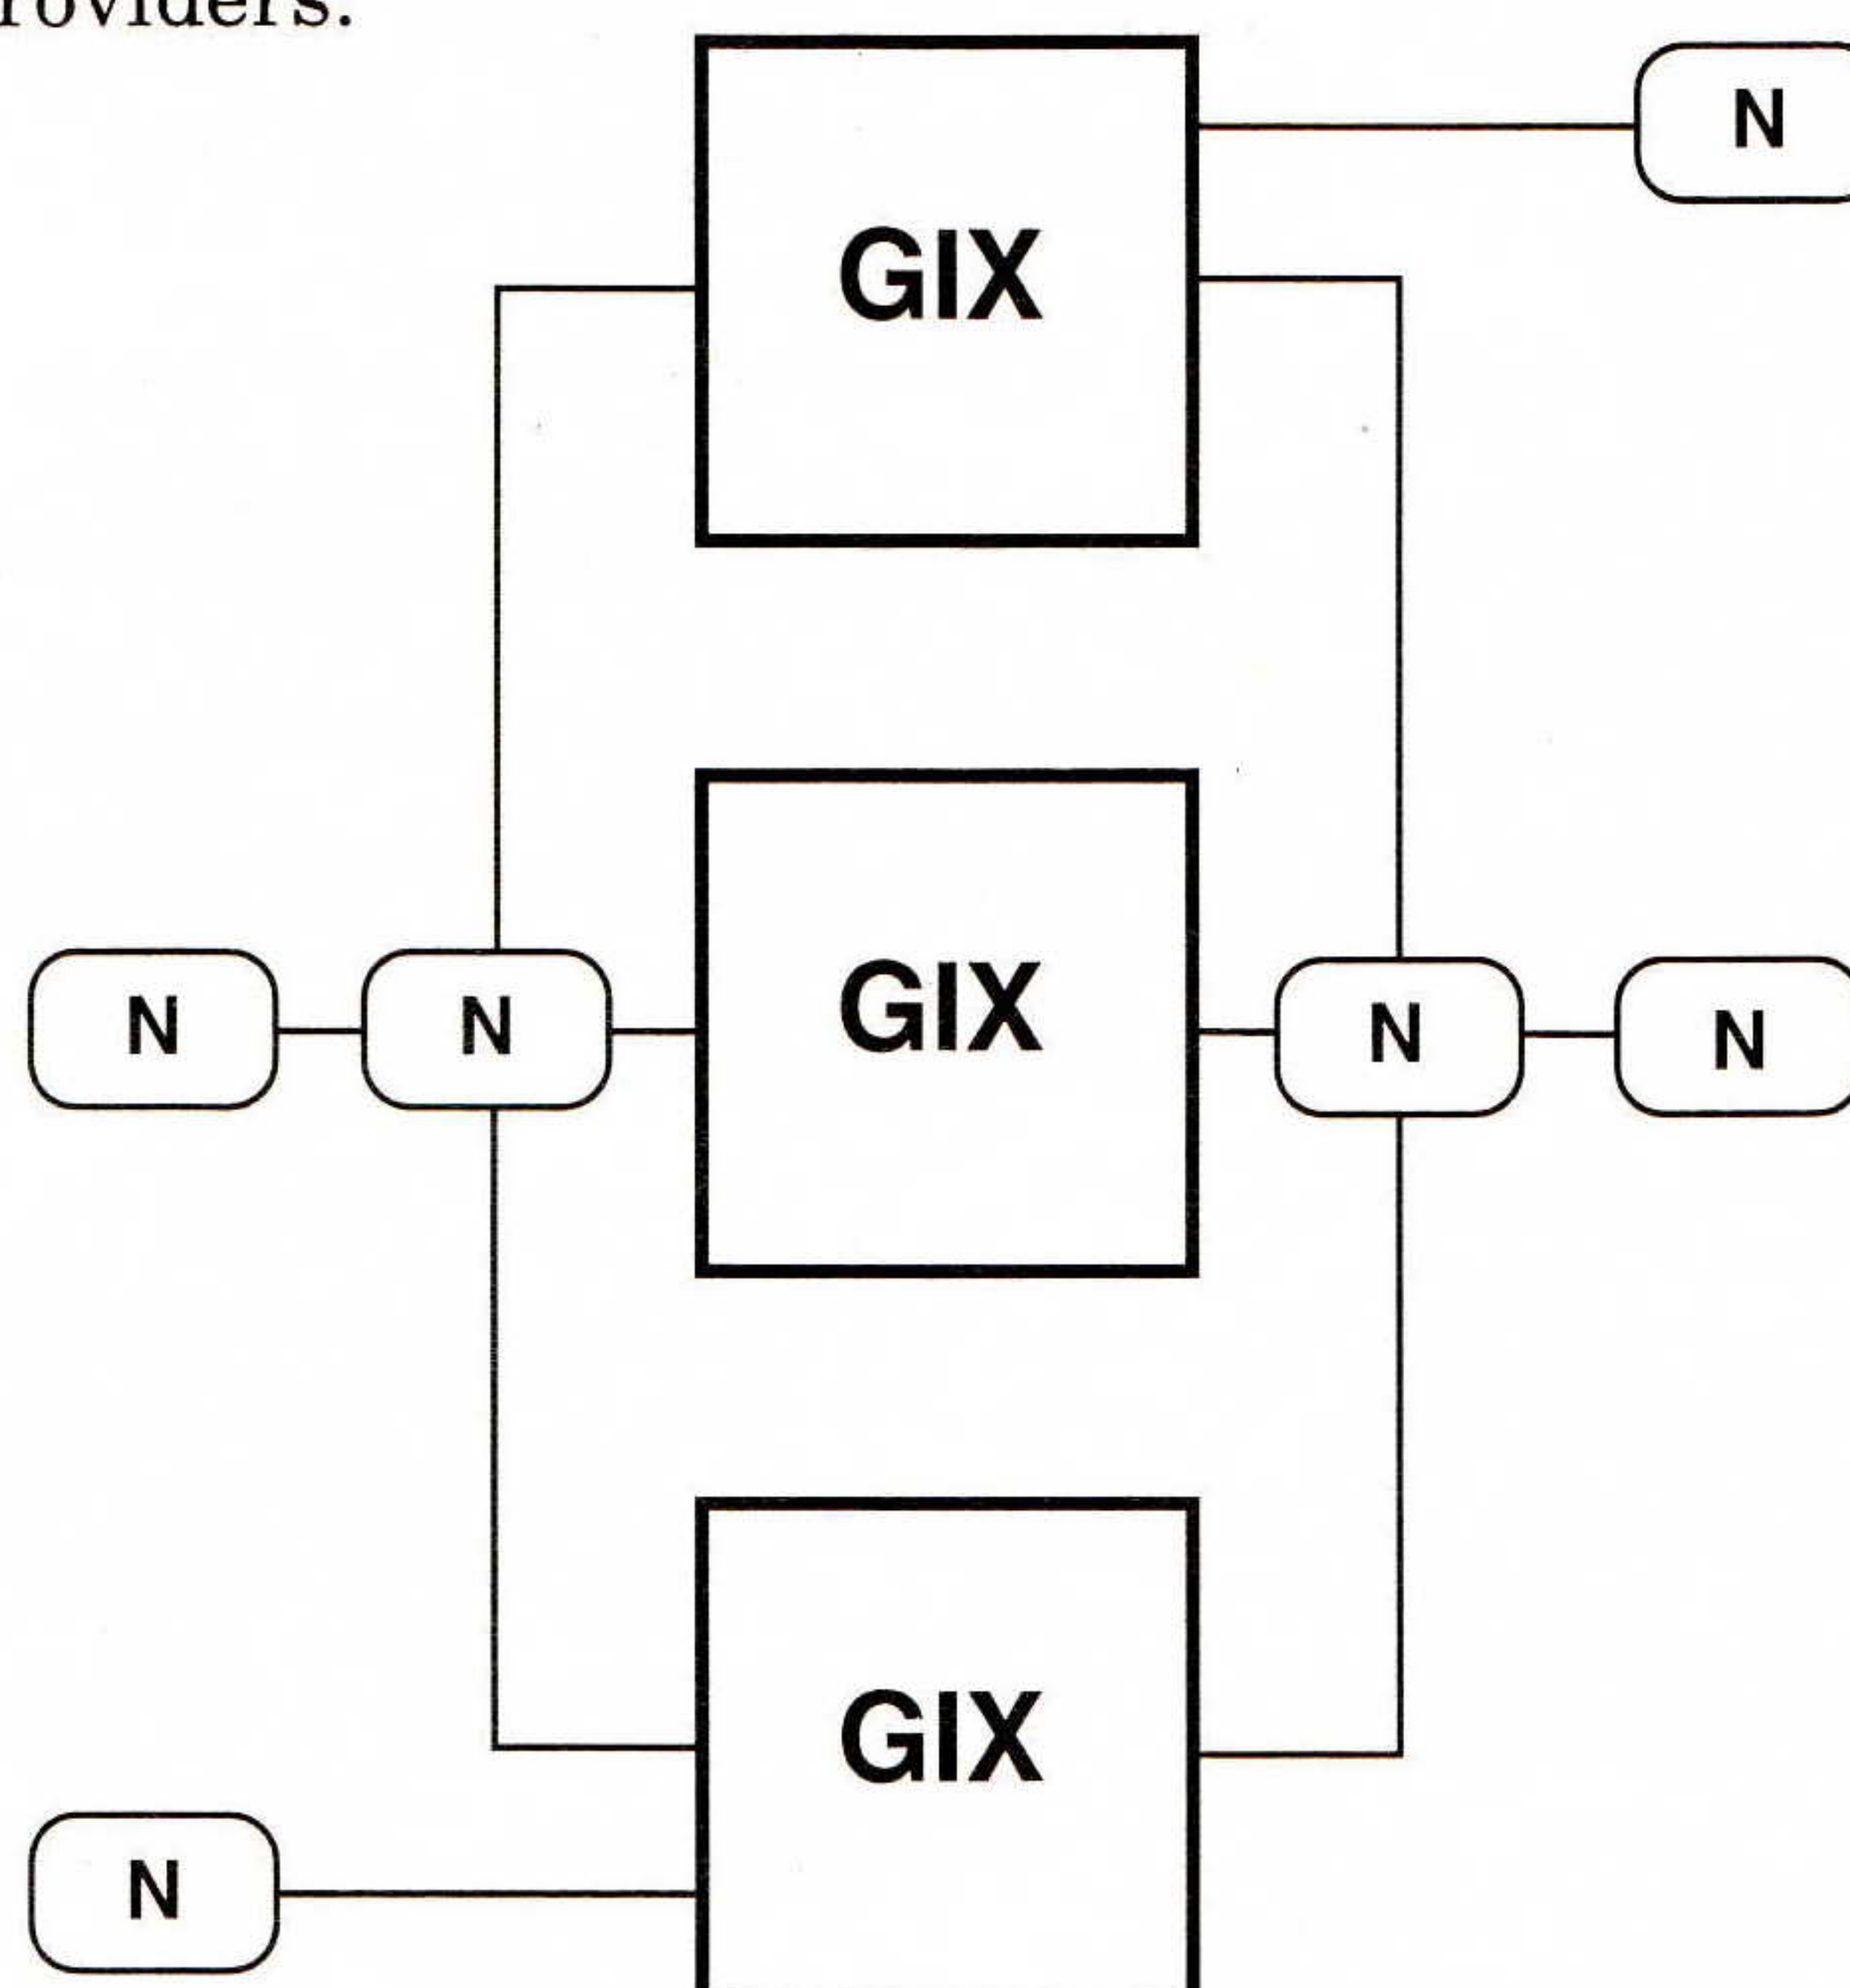
The basic GIX model with just one interconnection point where all the functionality is provided is shown below:



Networks located geographically close to such a GIX will have very low cost to connect as compared to networks being located on another continent.

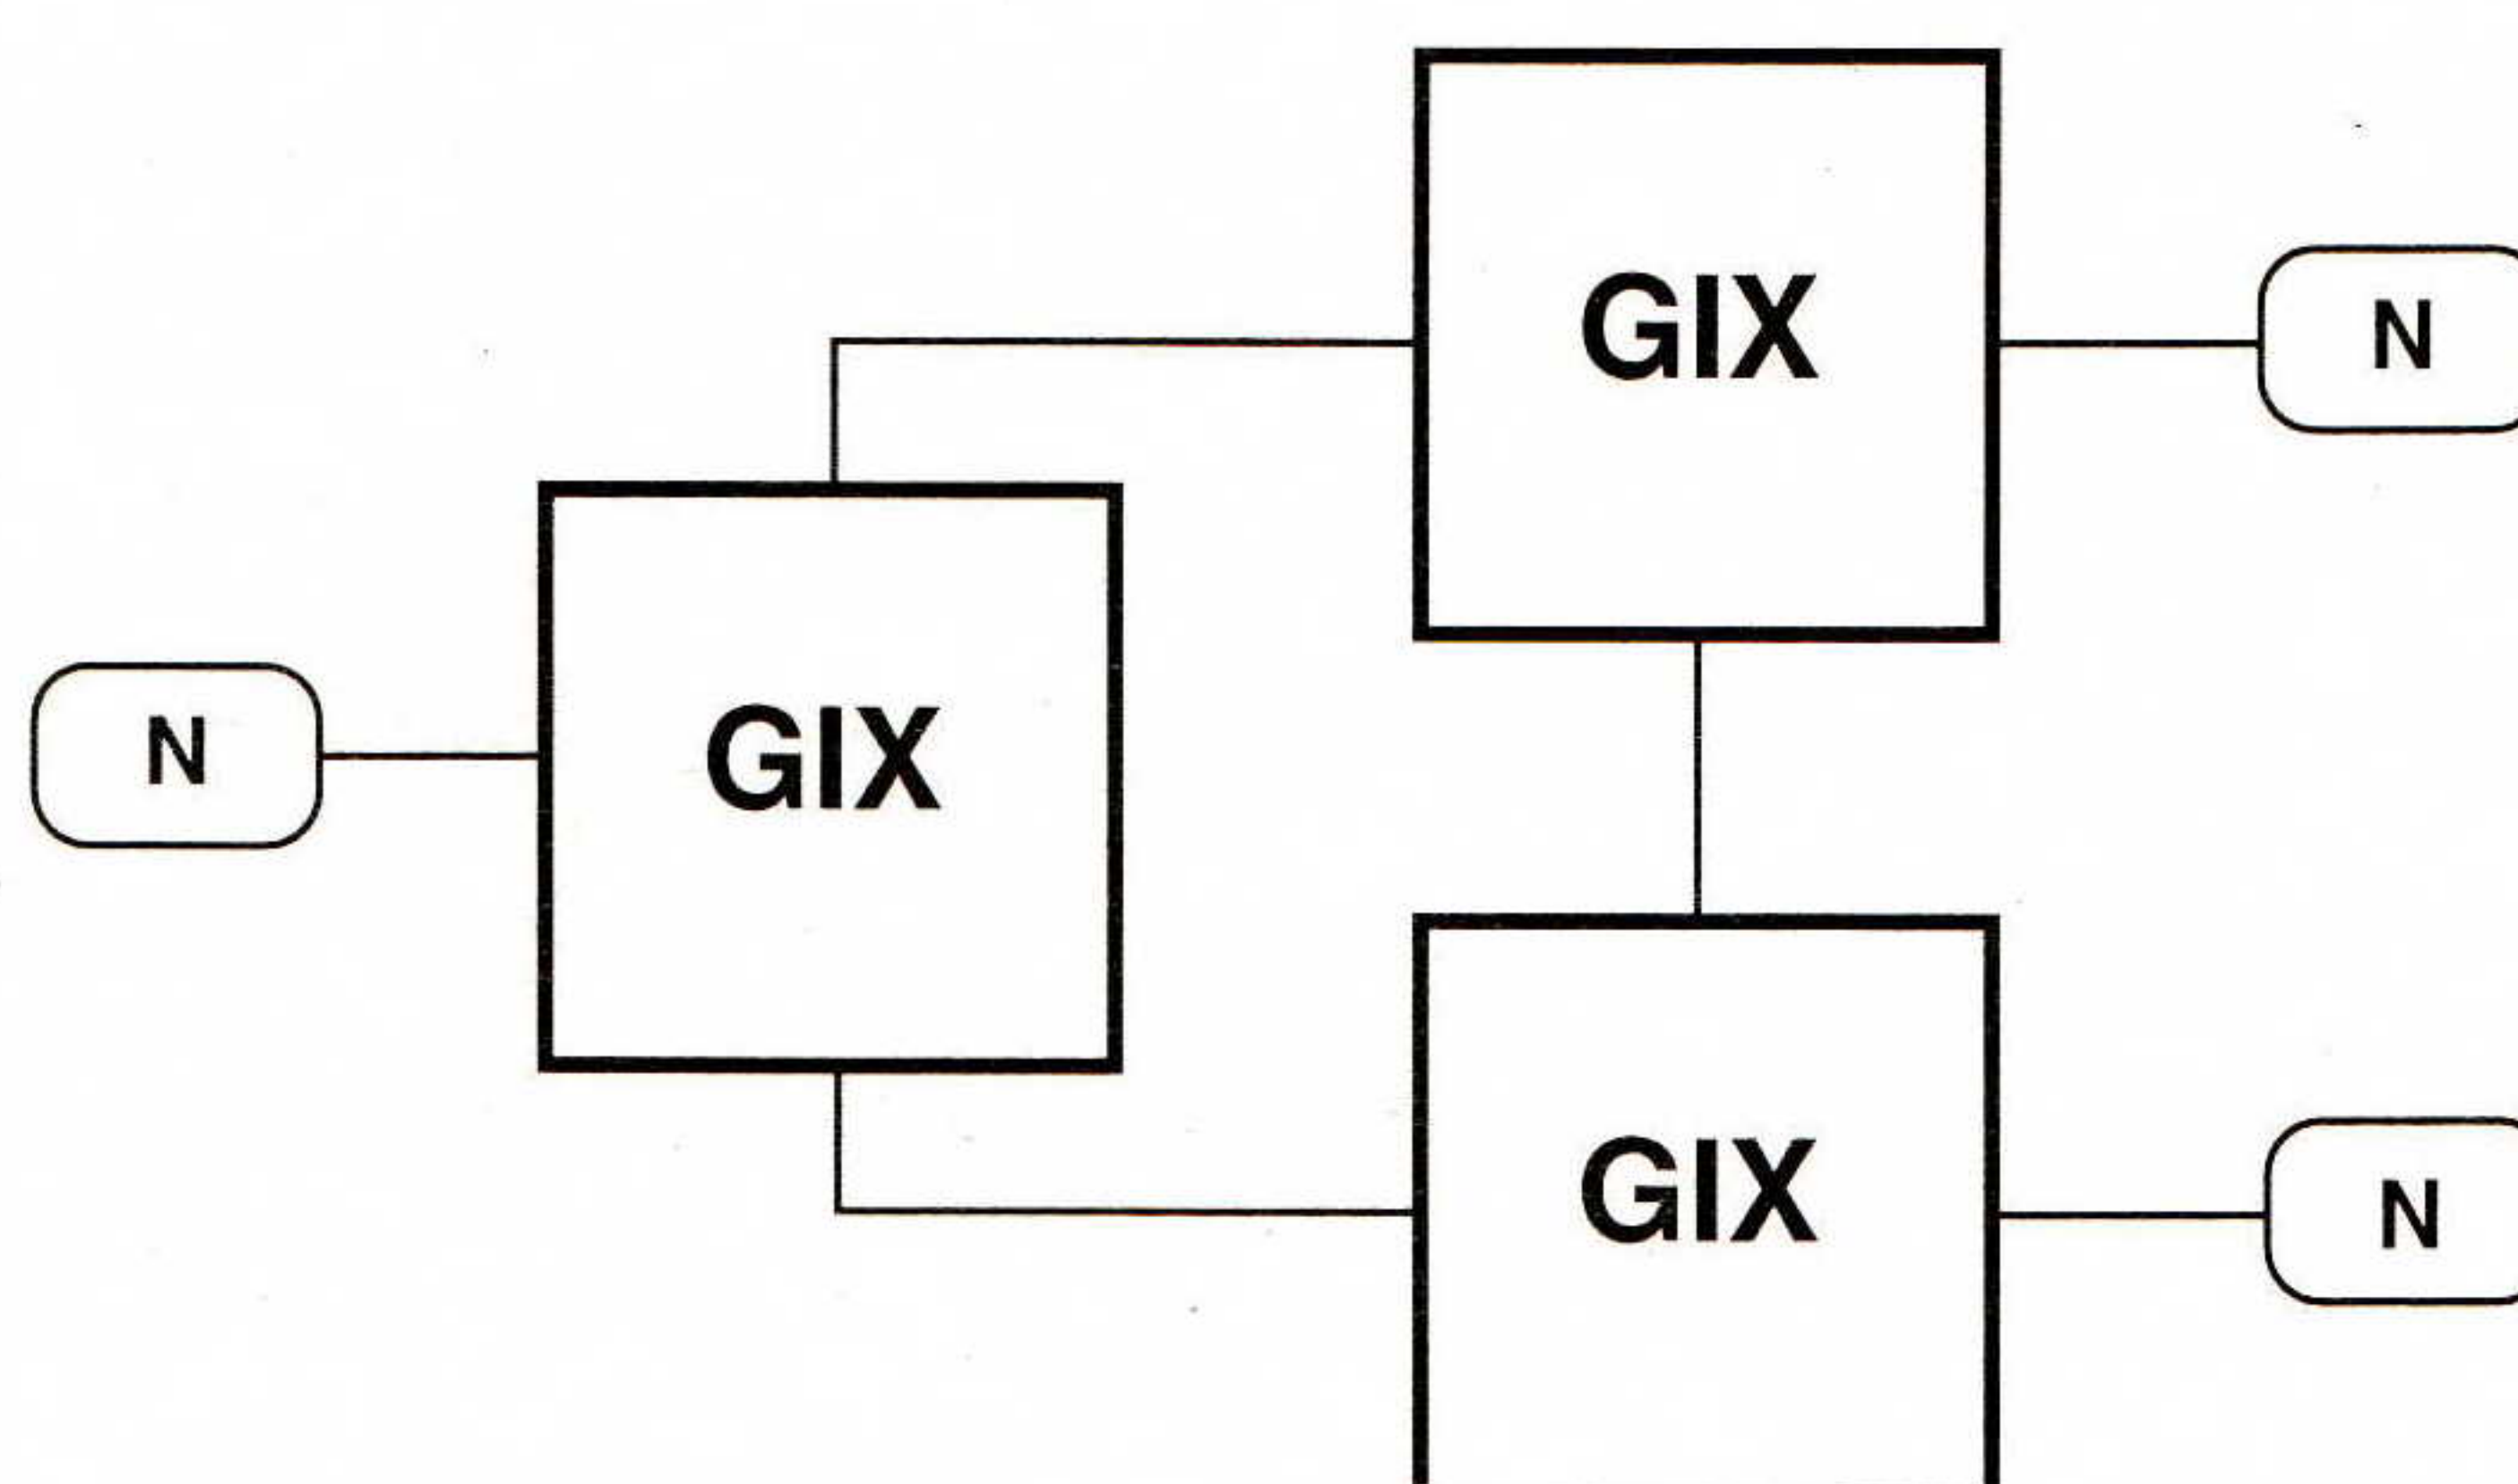
Global Internet Exchange (*continued*)

The next diagram shows multiple GIXs interconnected by several backbone providers:



In this scenario, some networks connect as backbone and transit networks to all GIXs. Other networks will have to negotiate transit capacity from these top level backbone networks.

Instead, it is possible for multiple GIXs to be interconnected by common shared resources as shown below:



Problems may arise when there is a need for upgrading the common shared resource. Finding agreeable methods for financing could prove difficult.

Refinements of the initial vision

At several major junctions in the global Internet, it is envisioned that there will be global interconnection points. Each such point would consist of a managed facility with 24-by-7 coverage and excellent environmental support. At each point there would be a high-speed broadcast LAN freely available for all kinds of traffic. Each participant would be free to, at its own expense, bring a circuit to this facility and place a router on the LAN. It would then be free to exchange routes and traffic with (the routers of) other participants at that LAN. Each participant would also pay for a pro rata share of the cost of the floor space, environmental support, and administrative support required for the interconnection points.

The Routing Registry

The task of the *Routing Registry* (RR) is to register preferred routes. This has to be performed on a neutral basis with respect to the different IP network providers.

Routing stability is implemented by having an RR register routes based upon the request from the network owner using similar procedures in place with the current NICs for network number registration. In the GIX context, the RR registers paths preferred from the GIX to the owner of the network.

The RR does not determine policy in any way. It just acts as a repository for routing information and performs housekeeping and consistency checking on the registered information. The result is a consistent view of the routing policies towards the ADs served by the RR and to a certain extent between those ADs as well.

It will be necessary to implement a method for clearly expressing the routing policy for any given IP network. This method should be easily understandable by today's network operators. One possible method is to describe routing policy in terms of Autonomous Systems (ASs).

The Route Server

The *Route Server* (RS) is pseudo-router running on a host directly connected to the physical GIX-LAN. The RS implements external routing protocols and maintains a routing table for destinations served by the RR operating the RS. The RS does not forward any packets but exports its routing table for use by the real routers on the GIX.

The task of the Route Servers is to disseminate consistent routing information to GIX connected routers. An RS first peers with all routers within its RR domain to acquire the dynamic routing information for ASs within this domain.

The RS then uses a "preferred path" database, derived from the routing registry of the RR to filter the dynamic routing information into a consistent routing table for its region. This routing table is then made available via BGP to all routers on the GIX that wish to use it.

It is anticipated that the number of routers at neutral interconnection points and the number of network numbers advertised by each router will both increase dramatically. In order to cope with this, it may be necessary to deploy a set of Route Servers on each of the interconnection points.

Networks that attach will have the option of using these Route Servers, of using traditional pair-wise peering, or of using some combination. There may, in fact, be multiple different Route Servers at a given interconnection point used by different sets of participants. The intention of Route Servers is not to impose policy, but to implement the dissemination of routes in a manner that scales and can be well managed.

A basic RS can be viewed as a virtual router making policy based routing decisions which it then communicates to the real routers. The real routers base their forwarding decisions on that routing information. The RS imports routes from all routers that connect to the domain that the corresponding RR has responsibility for, and exports routes to, anyone that starts a peering session with the RS. For redundancy and resilience it is assumed that eventually there will be at least two Route Servers per domain, peering with all AD routers and using IBGP between themselves to maintain consistency with each other.

Global Internet Exchange (*continued*)

GIX Pilot Project

Using Metropolitan Fiber System (MFS), a local carrier in some major US cities, Alternet, PSI, SURANet and Sprintlink have implemented a *Metropolitan Area Ethernet* (MAE) in Washington DC. The MAE, also known as "MAE-East," is based on MFS's 10Mbps Ethernet. MFS provides each connected site with an Ethernet AUI cable connection. To the sites the system appears to be an Ethernet LAN. MAE-East has been seen as a very good location for the first GIX pilot.

MAE-East is a broadcast medium. It could be hard to determine the source of bad packets appearing on it. In the long term someone will have to provide proper operations support. In the short term, the proposal should include procedures for handling faults.

By 1994 MFS should be able to provide connections at speeds higher than 10Mbps. Other communications suppliers have products similar to MFS.

It will be necessary to define a funding model for the GIX. A distributed interconnection could be a better way to connect routers to the GIX, since it minimizes the connection cost for each network provider. In the short term, the Washington MAE is the simplest way of providing connectivity to the Route Server.

The MAE-East can provide traffic capacity between providers. Pairs of networks needing greater traffic capacity may, however, make direct interconnects between themselves.

Rules

The working "rules" for connecting to MAE-East are as follows:

- Only network providers can join (no end users).
- Any network provider can join.
- MAE-East is "AUP-free" (AUP= *Appropriate Use Policies*)
- Any provider can peer (or not peer) with any other provider (mutual consent).
- Providers pay for "their" connection point.
- Should any costs be assessed for the GIX-East connectivity, they will be shared on a pro-rata basis (there are no costs being charged at this time.)
- No internal traffic is allowed on MAE-East.

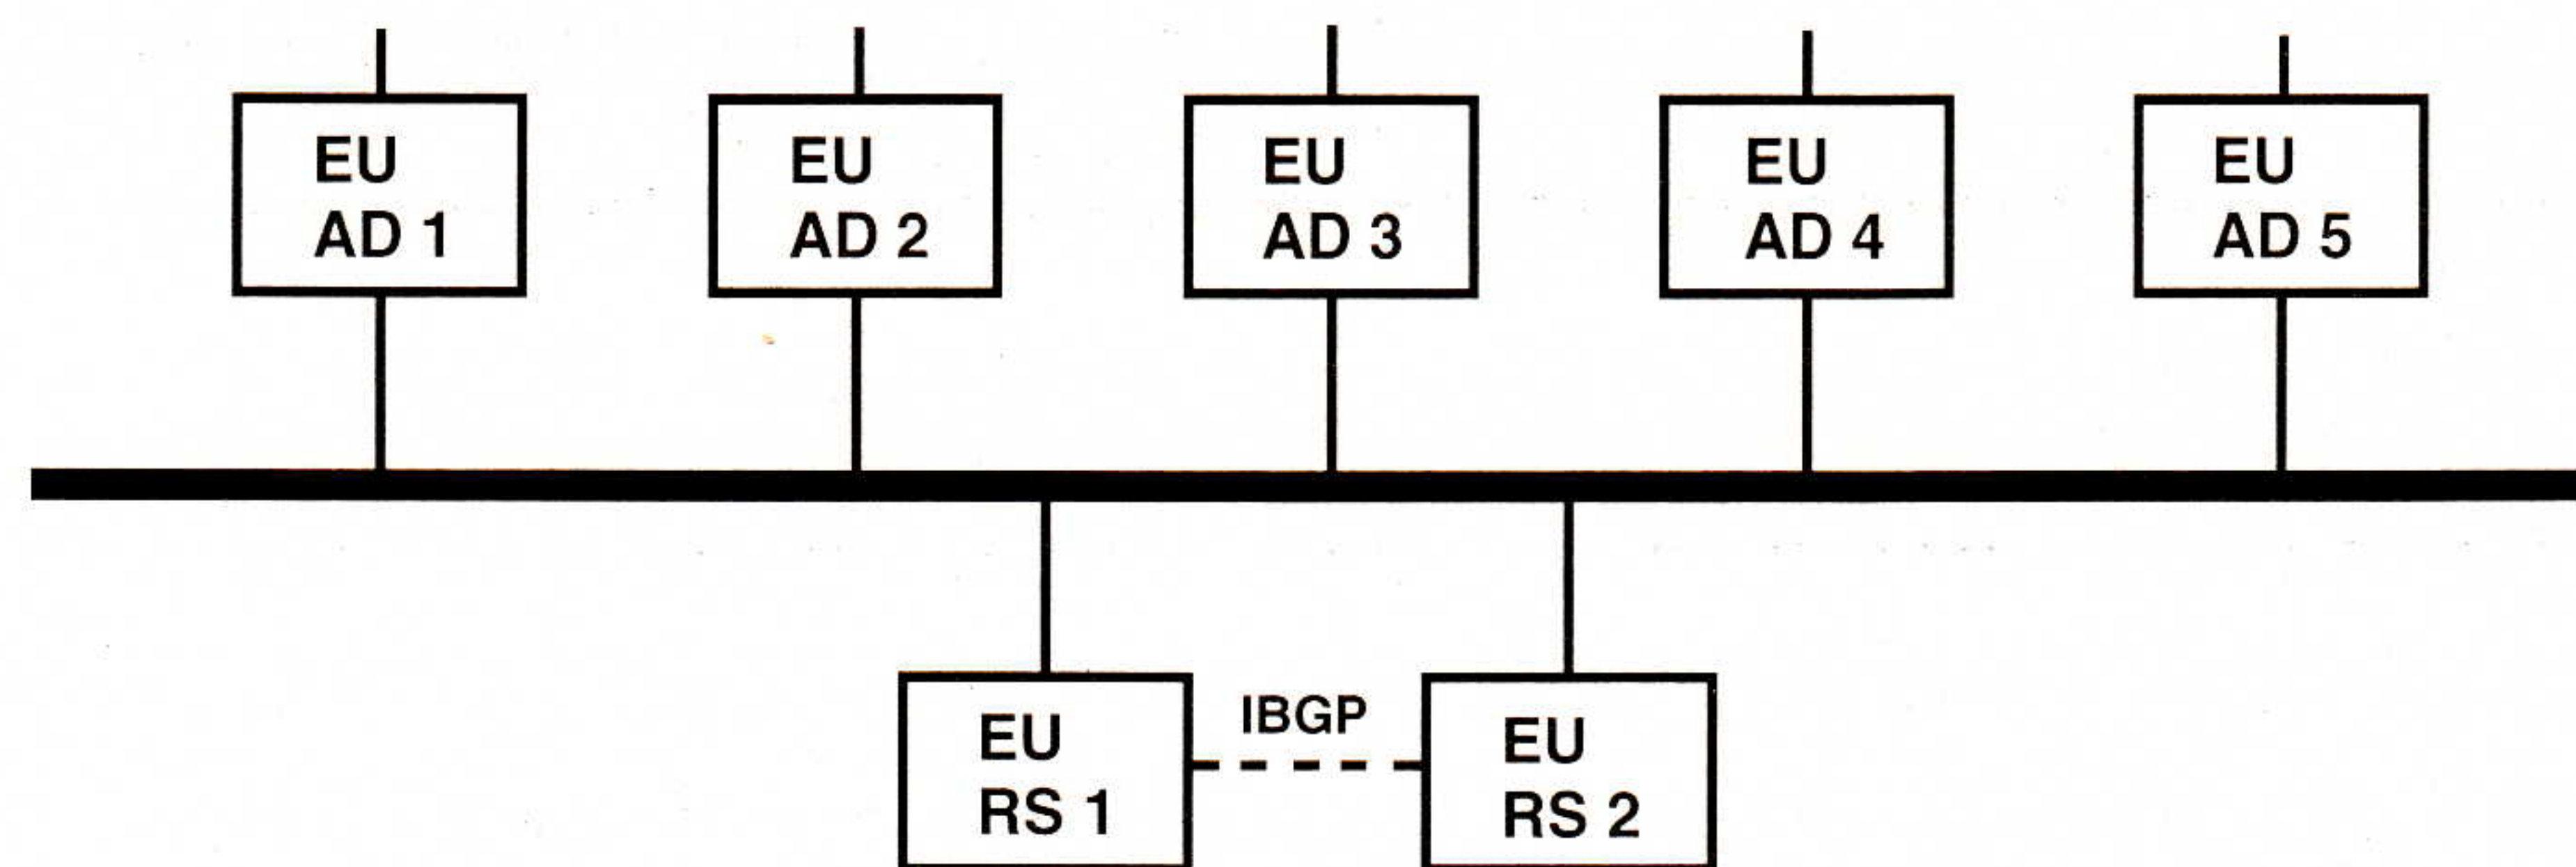
RIPE NCC Routing Registry

In Europe, the RIPE NCC has an effort underway to collect all the necessary routing information and store it in their database for European networks (see separate article in this issue). This is expected to be sufficient for the pilot. The current RIPE database has many of the essential objects needed to implement such a route filter mechanism today. Other regions will have to set up their own Routing Registry.

At the IETF meeting in Columbus, Ohio, in March 1993, two draft specifications of a policy description language were presented by Merit and by RIPE NCC. It is foreseen that these two proposals will be merged into one, and by that be submitted as an Internet standard.

European Route Server

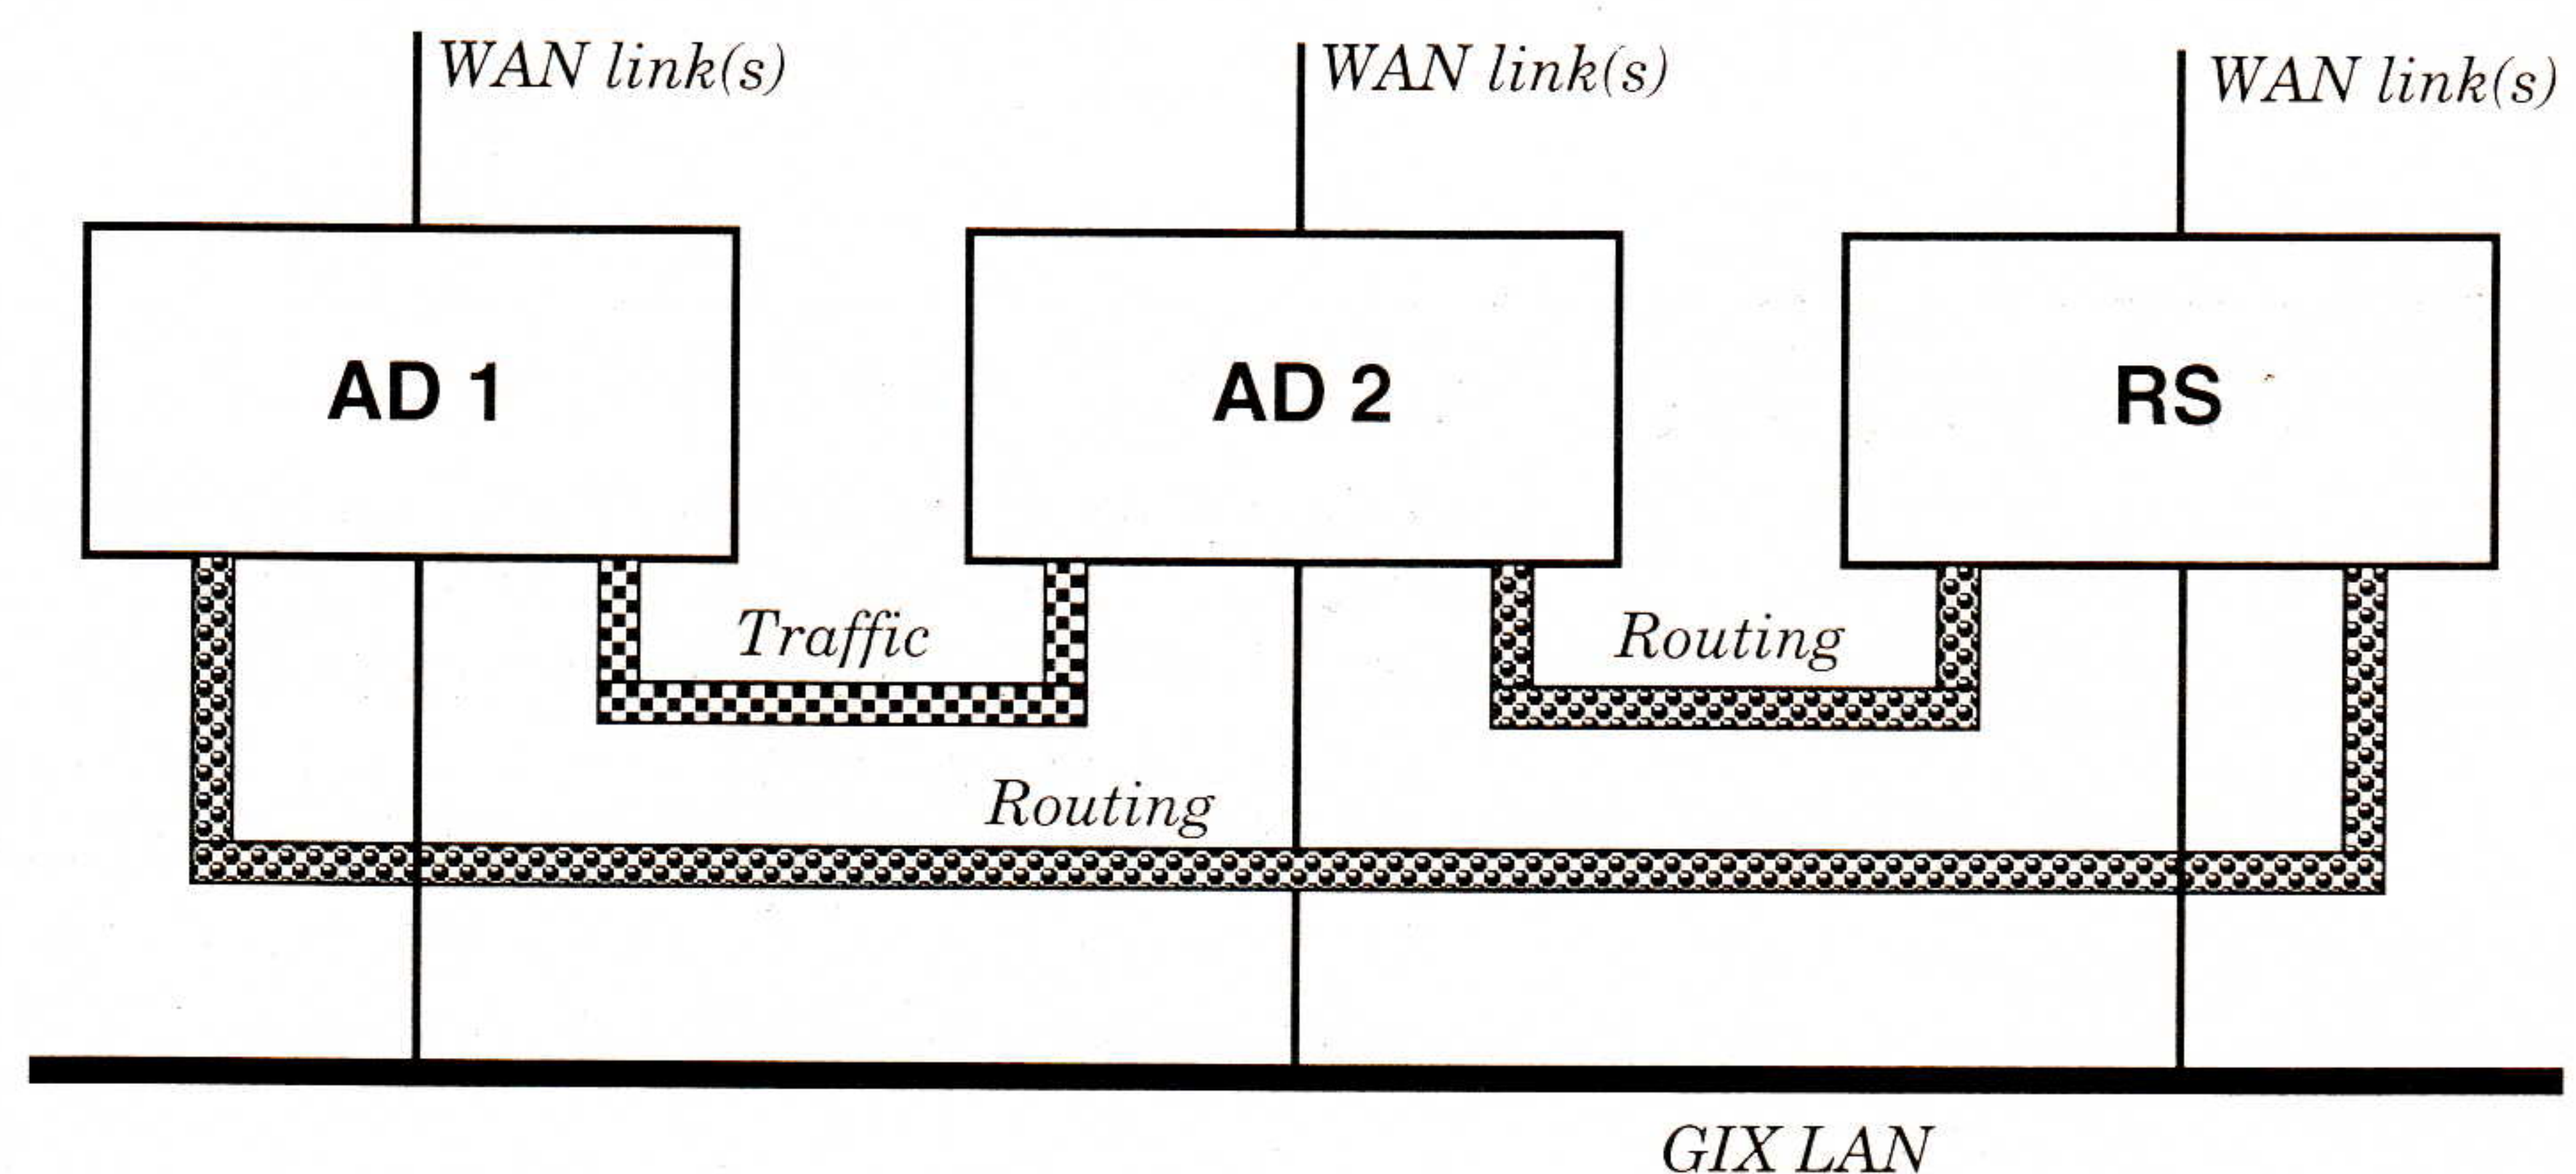
The figure on the next page shows the basic set up of the European Route Servers on the GIX. European routers EU-AD[1..5] are routers belonging to the different European ADs. Each European AD router peers using BGP with the two European Route Servers, EU-RS1 and EU-RS2. The Route Servers in turn peer with each other using IBGP.



The European Route Server will receive route announcements for all networks and filter them using the RIPE routing database. In addition, individual networks may pair across the LAN, which allows for exceptions to the general scheme.

The European Route Server will greatly improve the stability of routing between Europe and other continents. Questions of scaling will eventually have to be answered, but it is sensible to proceed with the pilot Route Server as quickly as possible so as to gather experience with it and to discover how it can be generalized according to the models described above.

It should be noted that when using BGP, the flow of routing information is separate from the general packet flow. This means that although routing information flows between the RS and connected AD-routers, the traffic will go directly between the AD routers. This way traffic flows across the GIX in only one hop as illustrated below:



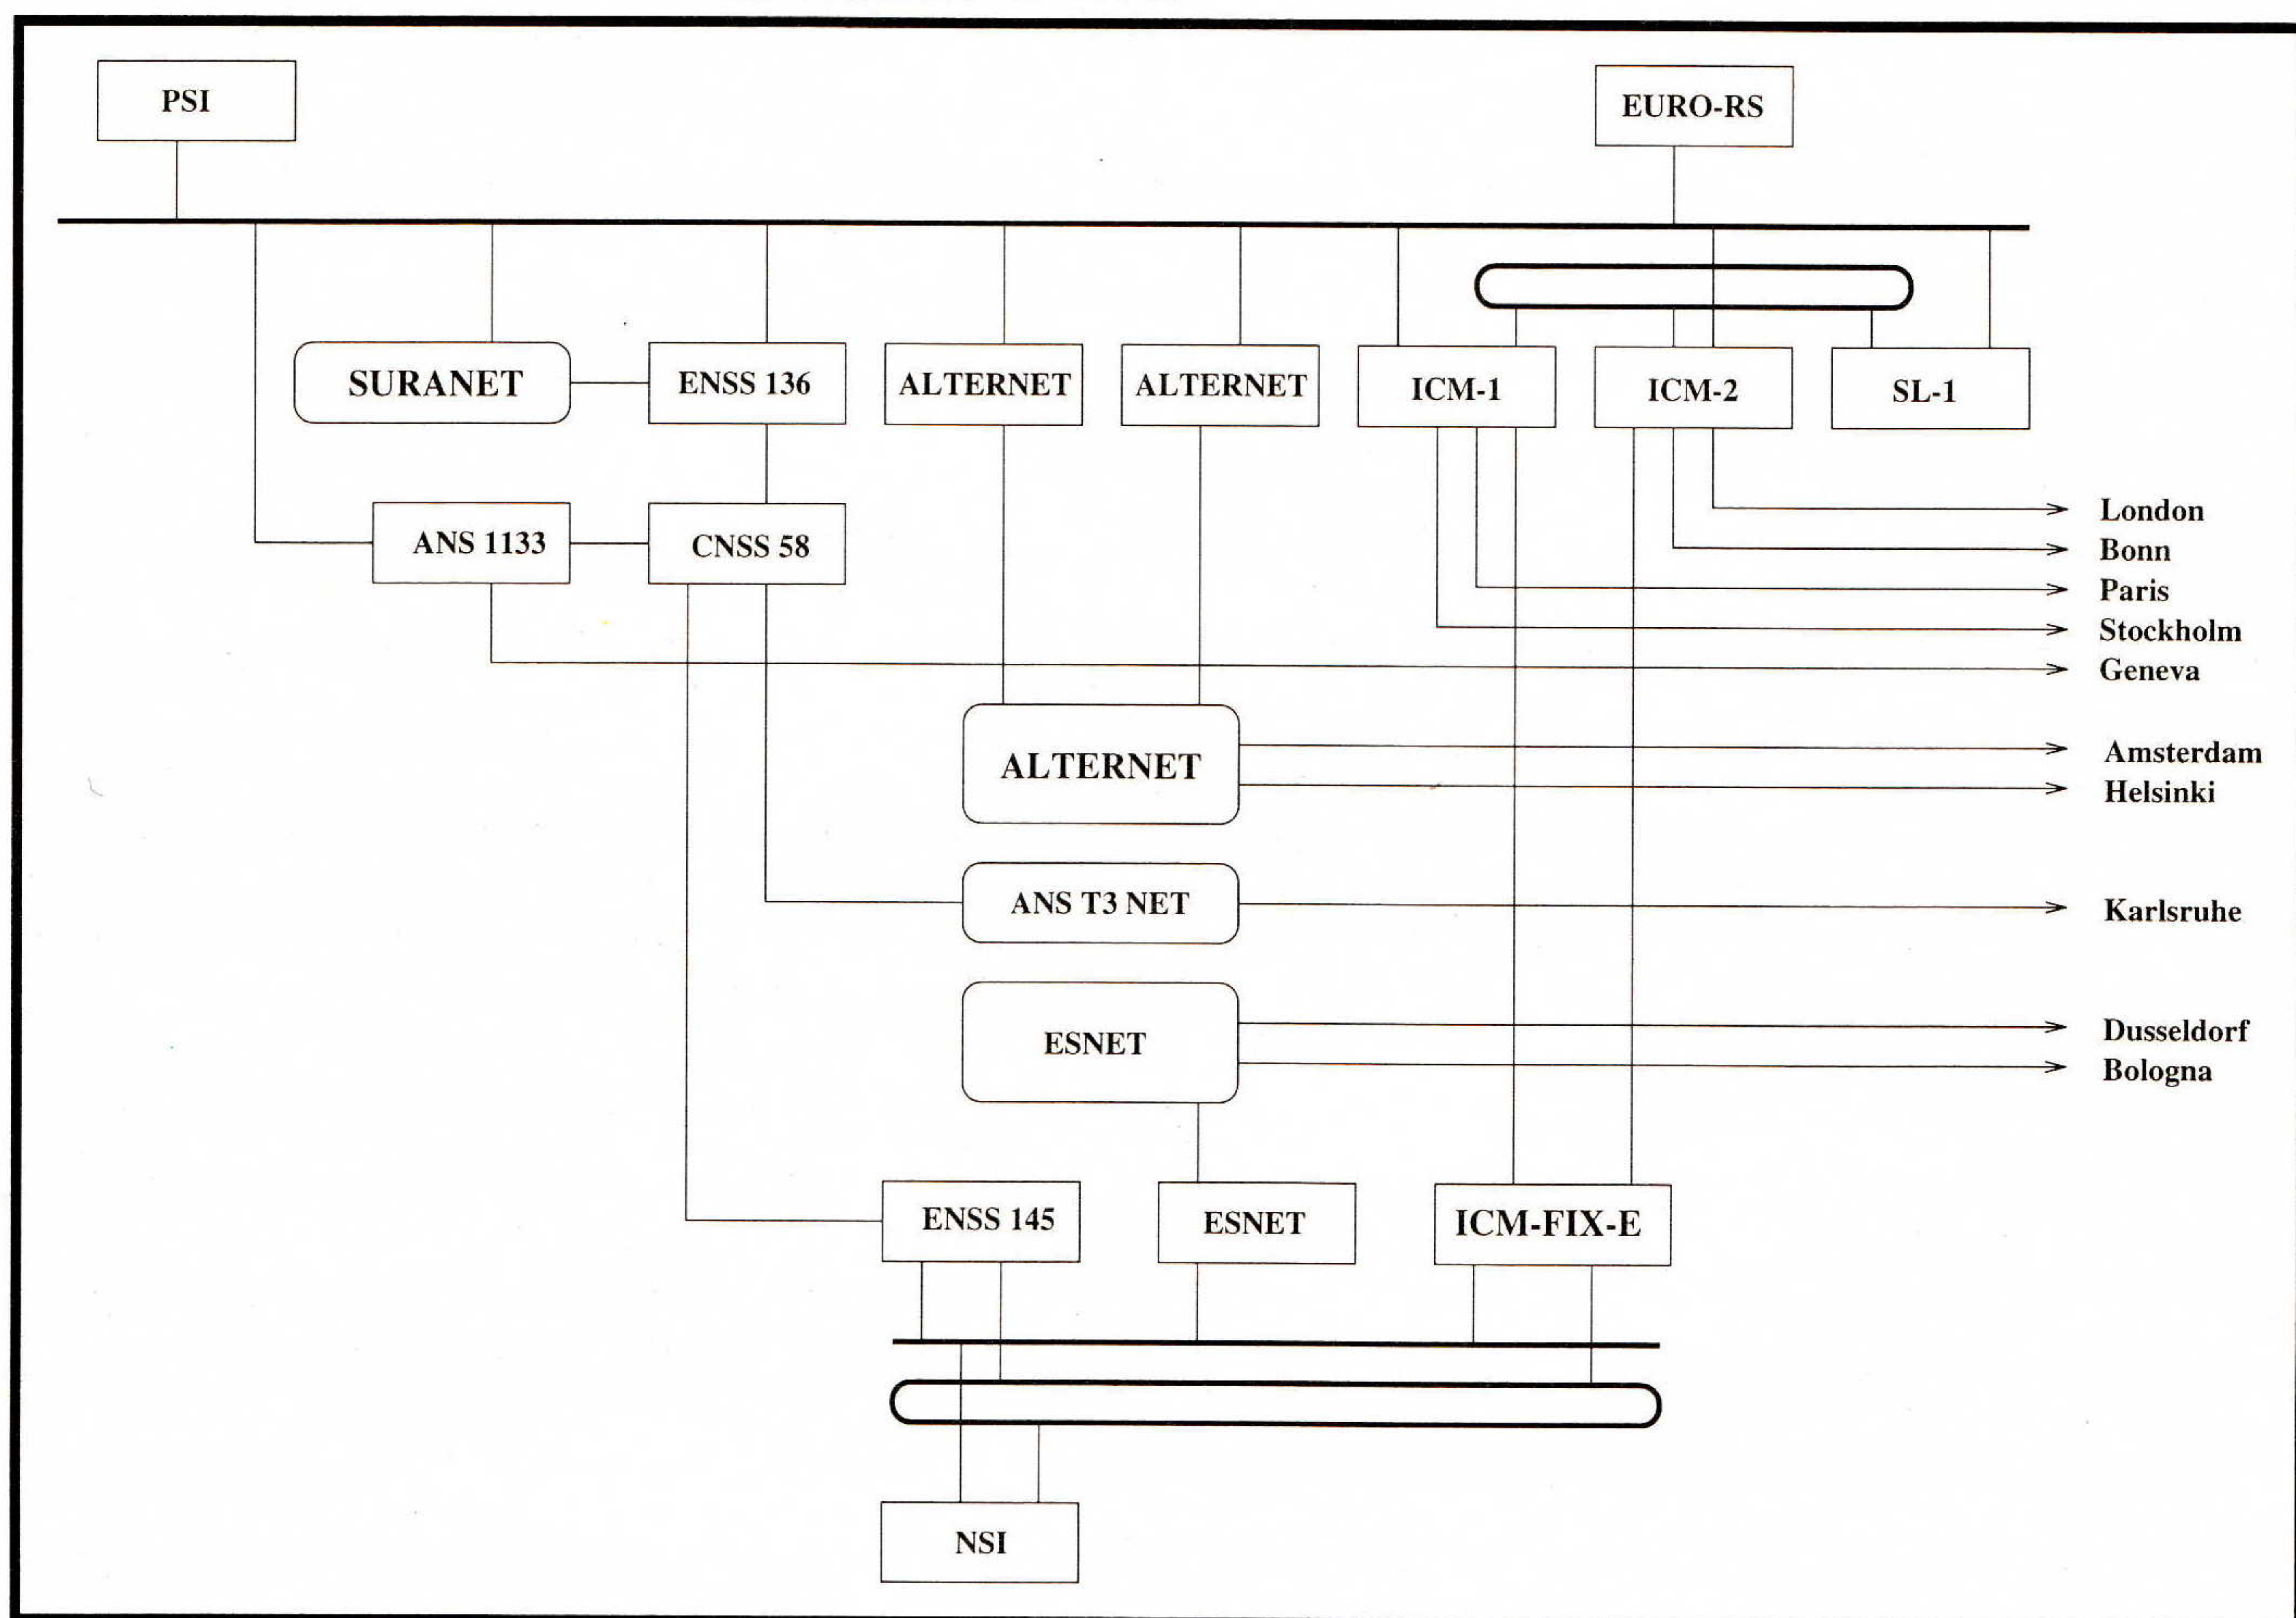
Furthermore, when there exists interconnections between ADs outside the GIX LAN it is fully possible that these ADs will use the GIX-LAN only to retrieve routing information. The data traffic will flow over private interconnections and never hit the GIX-LAN.

Other route servers at MAE-East

Merit has earlier announced interest in providing a Route Server for domains not covered by other Route Servers. At the IEPG/GIX meeting in Columbus, Ohio, March 1993 it was announced that the *Commercial Internet eXchange* (CIX) association will install a Route Server for its members. The European Route Server will be used as a basis. At this meeting it was also announced that Japan intends to provide similar service for the Asian and Pacific region.

Global Internet Exchange (continued)

The figure below shows the status of GIX connected networks as of the summer of 1993:



Future development

The current GIX will have problems in connecting all possible transit backbones due to differences in connections costs depending on the geographical locations. It is thus necessary to define a model for a distributed global interconnect system. Such system would have to be hierarchical as there is a need not only for global interconnects but also for regional and local interconnects. One challenge here is to define a fair model for distribution of costs among providers of inter-GIX connectivity between various GIX systems.

MAE-West

As a first temporary step to achieve a more distributed GIX connectivity, some efforts are ongoing in trying to extend the current Washington DC GIX to the US West coast ("MAE-West!") and to Europe using layer two technology. The intention here is to use existing "fat pipe" connections but split off a small fraction for routing purposes, and by that induce little or no cost for the transmission media. This is mainly a financial question, as the owners of the existing links have to agree on this as an acceptable way of using their part of the infrastructure. Another assumption is of course that the routing will only need an insignificant part of the current capacities (mostly T1/E1).

A method for the splitting of routing capacity has to be found. A "politically correct" solution might be the current "toy of fashion," *Asynchronous Transfer Mode* (ATM). Technically ATM does not, however, add functionality besides the ability of splitting of predefined amounts of bandwidth. ATM would thus act as a more flexible mux system, as compared to traditional TDM systems.

PRIDE

The RIPE NCC has for some time been urging European network organizations to describe their routing policy to be added to the RIPE routing database. At the moment, only around 50 percent of known networks have registered their policies. To define a financial and organizational framework for the ongoing efforts in developing tools and methods around the routing registry function, a project named PRIDE (*Policy-based Routing Implementation and Deployment in Europe*) has been initiated. See the RIPE information server on host `archive.ripe.net` for further details on this project.

Production mode

At the last IEPG meeting in San Francisco it was decided to move the current GIX pilot into production mode. Enhancements will still be considered but the basic functionality is now stable enough.

The RS could (and should) be enhanced to provide multiple routing tables or databases to different sets of AD routers. How to specify those features and how to guarantee overall consistency is beyond the scope of this article and is left for further study.

References

- [1] Guy Almes, Peter Ford, and Peter Lothberg, "Proposal for Global Internet Connectivity," June 1992.
- [2] Tony Bates, Daniel Karrenberg, Peter Lothberg, Bernhard Stockman and Marten Terpstra, "Internet Routing in a Multi Provider, Multi Path Open Environment," February 1993.
- [3] Tony Bates, Jean-Michel Jouanigot, Daniel Karrenberg, Peter Lothberg and Marten Terpstra, "Representation of IP Routing Policies in the RIPE Database," March 1993.
- [4] Daniel Karrenberg, "The RIPE NCC and the Routing Registry for Europe," *ConneXions*, Volume 7, No. 11, November 1993.
- [5] ATM Forum, "User Network Interface (UNI) Specification Version 3.0," August, 1993.
- [6] Heinänen, Juha, "Multiprotocol Encapsulation over ATM," RFC 1483, March, 1993.
- [7] Laubach, M., "ATM for your internet—But When?" *ConneXions*, Volume 7, No. 9, September 1993.
- [8] Atkinson, R., "Default IP MTU for use over ATM AAL5," work in progress, July 1993.
- [9] Laubach, M., "Classical IP and ARP over ATM," work in progress, July 1993.
- [10] Atkinson, R., "Towards Real ATM Interoperability," *ConneXions*, Volume 7, No. 8, August 1993.
- [11] Adams, A., "The Merit Policy-Routing Configuration System," *ConneXions*, Volume 7, No. 2, February 1993.
- [12] ed. Rekhter, Y., Li, T. "A Border Gateway Protocol 4 (BGP-4)," Internet Draft, December 1992.
- [13] Yakov Rekhter and Dave Katz, "The Border Gateway Protocol," *ConneXions*, Volume 5, No. 1, January 1991


*Learn more:
Session IT-5
Thursday at
10:30 am.*

BERNHARD STOCKMAN graduated as Master of Science in Electrical Engineering and Computer Systems from the Royal Institute of Technology in Stockholm, Sweden in 1986. After a couple of years as a researcher in distributed computer systems, he was employed by the NORDUNET and SUNET Network Operation Center. Today Bernhard is the European co-chair of the Intercontinental Engineering and Planning Group (IEPG) and involved in networking infrastructural questions both within Europe and on a global scale. He chairs the EBONE Action Team (EAT) responsible for the technical development of EBONE. His e-mail is: `boss@sUNET.se`

The RIPE NCC and the Routing Registry for Europe

by Daniel Karrenberg, RIPE

The European Internet



The RIPE NCC

Lacking the initial focus of ARPANET and the routing core function of the NSFNET, the European part of the Internet has grown to be a very diverse technical and organisational environment. A multitude of bilateral or multi-lateral interconnections between service providers has emerged. Even extremely valuable efforts such as EBONE have only partly managed to rationalise this topology and provide a measure of stability. The good connectivity and service enjoyed by European Internet users today is mainly due to the resilience of the basic Internet technology. The manageability and stability of the routing mesh is a matter of growing concern.

RIPE, the association of European IP service providers, has taken steps to tackle this problem. After shortly describing RIPE and the RIPE NCC, we will present the *European Routing Registry*. It is worth noting that concerns about routing stability are not unique to Europe. In the US, NSF plans to fund a *Routing Arbiter* function to deal with a very similar problem.

RIPE (*Réseaux IP Européens*) is a collaborative organisation open to all European Internet service providers. The objective of RIPE is to ensure the necessary administrative and technical coordination to allow the operation of a pan-European IP network. RIPE does *not* operate a network of its own.

RIPE has been functioning since 1989. Currently more than 60 organisations participate in the work. The result of the RIPE coordination effort is that the individual end-user is presented on their desktop with a uniform IP service, irrespective of the particular service provider his or her workstation is attached to. In August 1993, more than 450,000 hosts throughout Europe are reachable via networks coordinated by RIPE. The number of organisations connected to the European Internet is estimated to be more than 6,000.

The *RIPE Network Coordination Centre* (RIPE NCC) located in Amsterdam, The Netherlands supports RIPE and the European Internet service providers. It provides a wide range of technical and administrative support to network operators and the Internet community across Europe, including address space management, a *WHOIS* service, and an information service. Funding for this is provided by the service providers and other interested organisations. The NCC operates under the legal framework of the RARE association, the federation of European research networks.

In addition to the "core" support functions, the NCC provides a home to technical development projects which are useful to the Internet community. These projects are funded separately from the core activities by those interested in a particular project. One of these projects concerned with building the European Internet Routing Registry will be described later on, after a brief look at the NCC's core services.

The RIPE Database

In many operational situations ranging from loss of connectivity to intrusion attempts, it is necessary to establish quickly who is actually responsible for a specific IP network or DNS domain, and find a way to contact them. This information is available from the *RIPE Network Management Database* often referred to just as the "RIPE Database." Similar to the InterNIC's *WHOIS* database, the RIPE database contains information needed for the technical coordination of European IP networks.

Regional Internet Registry

It is maintained at the RIPE NCC in cooperation with local Internet Registries throughout Europe, and can be queried from anywhere on the Internet using a variety of methods including the standard WHOIS service. The information contained in the database will be exchanged with other regional NICs such as the InterNIC.

In order to operate IP networks in the Internet or between independent organisations, each physical network must use a unique IP network number. Other number spaces like the *Autonomous System* numbers need to also be maintained in a way that guarantees uniqueness. This is the task of the Internet Registry. Before the RIPE NCC existed these tasks were performed centrally for the whole Internet under the auspices of the US DoD and funded by various US agencies. Recently it has been recognised that the fair and hierarchical allocation of IP address space is a necessity for the Internet to survive.

To answer this need, the RIPE NCC has set up the *Regional Internet Registry* and a network of *Local Internet Registries* at IP service providers. The RIPE NCC has identified and helped set up more than 70 local registries to date. These local registries serve either the customers of a particular service provider or all organisations within a country which do not have a service provider. The latter category of registries is called *Non-Service Provider* registry, a community service which to date is being provided free of charge. 16 organisations willing to provide this have established a non-service provider registry.

Most European organisations now have access to a local registry which knows the local environment, speaks the local language, and can provide registration functions locally. Only requests for large amounts of address space are being reviewed by the NCC. As a result, IP number allocation to European organisations is now easier, more timely and, last but not least, *fairer* than before. The local support and the increased responsiveness are appreciated by network operators.

To date, the European Internet Registry system has allocated more than 10,000 network numbers, less than 1,000 of which are currently being routed on the Internet. This shows that many organisations are using the Internet protocols without actually connecting to the Internet, at least not initially. Another clearly visible trend is that the vast majority of new address space is assigned for purposes other than research and development.

Information services

In today's fast-changing Internet, everyone expects to find relevant information quickly and via the network. In order to support this, the RIPE NCC maintains the *RIPE Document Store*. The document store is a collection of RIPE's own documents and information relevant to the Internet in general such as RFCs, Internet Drafts and other technical documentation. Publicly available software tools which are useful for network operators are also made available via the document store. The document store can be accessed with traditional tools such as interactive login from the Internet, and via X.25 based networks and FTP. It can also be accessed with many of the new resource discovery tools such as *Archie*, *Wide Area Information Server* (WAIS), *Gopher* and *World Wide Web* (WWW).

Funding

Currently only a subset of European service providers contributes to the funding of NCC core services. This needs to change if the NCC is to be stable and thriving. All European Internet service providers need to contribute in a fair manner. Only in this way will the NCC be able to perform its core activities in a neutral and unbiased manner.

RIPE NCC and the Routing Registry (*continued*)

RIPE together with the current funders has developed a funding model to secure NCC funding in the future, and European Internet Service providers will have to commit to contribute their fair share to support the core activities of the NCC. Technical development projects will be funded separately from the core activities.

The Routing Registry

Within the membership of RIPE there is growing concern about routing stability and the management of an ever more complex Internet topology. The RIPE routing group has developed a routing stability strategy to deal with this in a very general way. The main premises of the strategy are: autonomy of network operators, independence of topology, and simplicity.

It has been clear from the start that any strategy that limits the autonomy of the network operators would not be accepted. This excluded any strategy dependent on a global routing core. Operators want to be free to establish routing exchanges with any other operator and have global stability at the same time. This meant that any strategy based only on a small set of exchange points (so called *DMZs*) would not work either. As a result, it became clear that the only way forward was to provide tools for the operators to manage their external routing better *themselves*. From this it followed that the solutions had to be simple and easy to understand in order to be deployed widely.

While the general routing stability strategy makes no assumptions about routing topology and supports a general routing mesh, RIPE still promotes the establishment of multilateral exchange points and a somewhat hierarchical routing structure including a GIX (*Global Internet eXchange*—see elsewhere in this issue). This will undoubtedly make the problem less complex and make it easier to achieve stability.

The general strategy, however, is to provide tools for the network operators to simulate, configure and troubleshoot their external routing and diagnose connectivity problems. These tools need a lot of non-local information to function. Because external routing always depends on more than one operator, a means to exchange information about the respective routing configurations and policies is necessary. In order to diagnose distant connectivity problems one has to know the policies of distant networks. All this information will be maintained in the *Routing Registry*, a database operated on neutral ground by the RIPE NCC.

Format

After thorough discussion and experiences with a previous format, in the spring of 1993 RIPE established consensus on a format to describe external routing policy and other necessary information. The descriptions can be registered in the routing registry, making them available for perusal and accessible to the tools. This makes it possible to check consistency of the descriptions with the descriptions of neighboring operators, diagnose connectivity problems and more. The format is described in document **ripe-81**.

The basic idea is to describe *local* external routing exchanges on a very low level, not much above the level at which the routers are configured. In this way, only the direct exchanges between neighboring network operators are being described. It is these exchanges that the network operators have direct authority over and that the operations personnel are familiar with. The low level description also prevents the specification of policies which cannot be implemented by today's technology.

The Whole is more than the Sum of the Parts

From these descriptions, actual router configurations can be derived in a rather straightforward way. At least the network operators can compare the configurations generated with their current configuration in order to detect discrepancies. Eventually, after enough confidence is built, the generated configurations can actually be used to configure the routers. This ensures that the descriptions in the Routing Registry actually reflect the real configuration.

By collecting all these locally generated and maintained descriptions, a picture of the whole European routing mesh is created. This enables diagnostic tools to describe the connectivity between service providers, and more importantly find areas with no connectivity. Another tool will combine the information about the actual path traffic takes as reported by *traceroute* with information in the registry. From the combined information it will be obvious immediately whether traffic follows and intended route and if it is the primary or a fall back route. A prototype produced the following output:

```
prtraceroute jolly.nis.garr.it
```

```

1 AS1104 hef-router.nikhef.nl      (192.87.45.80)  [I]
2 AS1755 Amsterdam-EBS1.Ebone.NET (192.87.4.17)   [D1]
3 AS1755 Cern-EBS1.Ebone.NET      (192.87.4.10)   [I]
4 AS 513 chep1.cern.ch            (192.65.185.2)  [D1]
5 AS 137 GARR-gw-mi.infn.it       (192.12.193.49) [E1]
6 AS 137 GARR-gw.cilea.it         (192.12.193.41) [I]
7 AS 137 GARR-gw.cineca.it        (192.12.193.21) [I]
8 AS 137 GARR-gw.cnr.it           (192.12.193.14) [I]
9 AS 137 jolly.nis.garr.it        (192.12.192.5)  [I]
```

```
Path followed: AS1104 AS1755 AS513 AS137
```

```

AS1104 = NIKHEF
AS1755 = EBONE-INTERNAL
AS 513 = CERN-AS
AS 137 = GARR/INFN IT
```

The second column of the trace describes the autonomous systems traversed (see the legend at the bottom of the trace). The last column gives information about the routing policy used for each hop. Internal hops within an operator's Autonomous Systems are labeled **I**, hops between ASs are labeled **D** or **E** depending on whether a default route is being used. The numbers associated with external hops describe whether a primary (1) or a fallback route (2, 3, ...) is being followed. If a path is observed that does not conform to the registered policies this will be indicated by question marks. This presentation of the data makes it immediately obvious if something is wrong and where. Operational personnel need not have external knowledge about routing policy. The diagnostic information is presented in combination with the policy information.

Population of the Registry is key

It is evident that all this can only be useful if all network operators register their routing policies in the routing registry and use the registry based tools. Currently only about half of all European autonomous systems have documented (part of) their external routing exchanges in the Routing Registry. In order to improve this, the PRIDE project (*Policy-based Routing Implementation and Deployment in Europe*) has been started at the RIPE NCC. In this project the tools will be implemented and training will be offered to the network operators in order to motivate and enable them to make use of the routing registry.

continued on next page

RIPE NCC and the Routing Registry (*continued*)

PRIDE is funded by the Norwegian Research network UNINETT, the UK Joint Network Team, Unet Communications Services Inc. and EUnet B.V.

Conclusion

RIPE has developed a viable strategy to ensure routing stability in the ever growing European Internet which can deal with an ever more complex routing mesh.

Author's Address:

Daniel Karrenberg
RIPE NCC
Kruislaan 409
NL-1098 SJ Amsterdam
The Netherlands
Tel: +31 20 592 5065
Fax: +31 20 592 5090
E-mail: ncc@ripe.net

More information can be obtained via anonymous FTP to Internet host `ftp.ripe.net`, by telnetting to `info.ripe.net`, or via *Gopher* to `gopher.ripe.net`.

References

- [1] Stockman, B., "EBONE, The European Internet Backbone," *ConneXions*, Volume 7, No. 5, May 1993.
- [2] *ConneXions*, Volume 7, No. 5, May 1993, "Special Issue: Focus on Europe."
- [3] Guy Almes, Peter Ford, and Peter Lothberg, "Proposal for Global Internet Connectivity," June 1992.
- [4] Tony Bates, Daniel Karrenberg, Peter Lothberg, Bernhard Stockman and Marten Terpstra, "Internet Routing in a Multi Provider, Multi Path Open Environment," February 1993.
- [5] Tony Bates, Jean-Michel Jouanigot, Daniel Karrenberg, Peter Lothberg and Marten Terpstra, "Representation of IP Routing Policies in the RIPE Database," March 1993.
- [6] Deutsch, P. & Emtage, A., "The *archie* System: An Internet Electronic Directory Service," *ConneXions*, Volume 6, No. 2, February 1992.
- [7] Kahle, B., "An Information System for Corporate Users: Wide Area Information Servers," *ConneXions*, Volume 5, No. 11, November 1991.
- [8] McCahill, M., "The Internet Gopher: A Distributed Server Information System," *ConneXions*, Volume 6, No. 7, July 1992.
- [9] Berners-Lee, T., "A Summary of the WorldWideWeb System," *ConneXions*, Volume 6, No. 7, July 1992.
- [10] Stockman, B., "Global Connectivity: The Global Internet Exchange (GIX)" *ConneXions*, Volume 7, No. 11, November 1993.

*Learn more:
Session IT-5
Thursday at
10:30 am.*

DANIEL KARRENBERG has been active in European networking since 1981, first at Dortmund University, from where he received a graduate degree in computer science in 1987 and later at CWI in Amsterdam, Netherlands. Daniel has helped to create the pan-European network EUnet which at first offered UUCP services has since become a full Internet service provider. He is currently responsible for the RIPE Network Coordination Centre which coordinates the European part of the Internet and acts as the European regional Internet registry.

Profile: EUnet

by Glenn Kowack, EUnet

European networking priorities

Discussions within the European Internetworking community are often led by the government-sponsored national and international research networks. Those discussions often focus on the quest for "big bandwidth" and the massive governmental support, funding, and intervention that it demands. However, well before these discussions became prominent, a different model of infrastructure building was at play, steadily creating a large fraction of the greater European Internet. The most extensive and longest-operating European champion of this different model is EUnet.

This article

This article overviews EUnet and its demonstration of an alternate and highly successful route toward the development of the Internet and, most importantly, in how it serves real customers (not just the perennially invoked, yet somehow never-present "users"). It traces the history of EUnet, and our philosophies of service provision and network economics. It is not intended to be an indictment of the desire or hard work done by many in the quest for big-bandwidth academic networking before the marketplace will support it. Rather, it points to an approach which has produced significant end-user service and will certainly bring us to substantially greater bandwidth levels in the near future (if American experience is any guide, it will not be more than a few years away).

EUnet snapshot

EUnet began in 1982 as a UUCP based network concentrating on electronic mail. EUnet has since grown to become perhaps Europe's most extensive Internet-related service organizations, serving 10,000 sites and networks via 50 points of presence. With service centers in 27 countries providing coverage from Iceland to Vladivostok and from the Arctic Circle to Northern Africa, EUnet has one of the widest geographical spans of any Internet provider in the world. Traffic is carried nationally, internationally, and intercontinentally over EUnet's own infrastructure. All of this has been accomplished without direct major governmental sponsorship or funding.

EUnet core services include e-mail (RFC 822 and X.400), File Transfer (FTP), Remote Login (Telnet), Network News, and Archive Access. Connectivity is available via leased, X.25, ISDN, and analog dial-up connections, using IP and UUCP protocols.

EUnet's National Service Providers include:

Austria	Belgium
Bulgaria	Czech Republic
Denmark	Finland
France	Germany
Great Britain	Greece
Iceland	Ireland
Italy	Luxembourg
The Netherlands	Norway
Portugal	Slovakia
Slovenia	Spain
Switzerland	The former Soviet Union
Tunisia	

EUnet Providers are being set up in:

Algeria	Egypt
Poland	Romania

Profile: EUnet (*continued*)

Early history

EUnet was announced by members of the European computer industry and researchers from institutes and universities at the April 1982 meeting of the *European Unix Users Group* (EUUG—in 1990 the EUUG was renamed “EurOpen”). At that time, there were no generally-available means for researchers or anyone else to exchange electronic mail either internationally within Europe or across the Atlantic to the United States. In most European countries, there were no national research networking organizations (notable exceptions were JANET in Great Britain, and HEPnet), and most institute-to-institute communications were provided in an *ad hoc* fashion, if at all.

The UUCP and USENET developments in the United States had by 1982 already shown that a rather anarchistic, low-overhead, yet effective system of networking was possible without major government or institutional support. However, the differences among the countries of Europe meant that the simple “find a willing host and exchange traffic” model of the USENET could not be transported wholesale to Europe; the much higher European telephone tariffs, the frequent telecom border crossings, and differences in national languages, technical development, and regulations required a different solution.

The approach chosen revolved around the designation of a single point of concentration for e-mail exchange within each participating country. Each national concentration point would then forward all international mail to a single European concentration point. This two layer hub-and-spoke arrangement proved to be the perfect European interpretation of the original USENET approach. Groups participating in the early days of EUnet included the Institut National de Recherche en Informatique et Automatique (INRIA) in Paris, the Royal Institute of Technology in Stockholm (KTH), the University of Kent at Canterbury, the University of Copenhagen, and the University of Genoa. The Center for Mathematics and Informatics (CWI) in Amsterdam acted as the European hub and network operations center, providing connectivity within Europe and to the United States.

EUnet transport was originally based upon UUCP, and services were primarily RFC 822 e-mail and Network News. Within a few years, nearly all of the European Community region was covered, which provided enough income and traffic so that by 1986 our connection to the US was upgraded to a leased line running SLIP. In early 1987, leased lines were in use between the European and national NOCs; later that year leased line IP services were provided directly to subscribers (today marketed under the name InterEUnet).

Access to EUnet services was originally restricted to members of the research and development community, again following the American model (this time unfortunately). This was mandated by monopolistic telecom regulations in the European countries, and reinforced by the general pre-commercial practices of the day. This constraint no longer applies, as we will see below. Each national EUnet networking team worked in conjunction with a national EUUG group. The EUUG overall provided support, cohesiveness, and a common forum for EUnet development.

National EUnet evolution

Nearly all National EUnets were begun by networking enthusiasts at universities or national research institutes. Starting with a few workstations and modems, most began by providing simple e-mail and Network News services to academic institutions, and later to research groups in industry.

Each subscriber was required to pay for the networking services they used. As the number of subscribers grew, it eventually became possible to obtain a leased line to the Amsterdam NOC, thus enabling national distribution of the full set of Internet services. Steady evolution permitted EUnet staff to develop an in-depth understanding of the economics and technology of service provision. It also continues to be a very effective start-up model for use in countries which, now just beginning to join into the integrated economy of Europe, have serious national telecom infrastructure problems to overcome. And, the "pay actual costs of use" approach imposed economic discipline, still sorely lacking in many government-sponsored nets, from the subscriber through the entire EUnet organization. EUnet subscribers have over ten years of experience in understanding how networking fits into their overall cost picture; there are no mysteries about actual cost and its relationship to the value of network use. EUnet has proven that end-use of networks is economically viable and desirable, even given the artificially high European telecom prices and the difficult technical circumstances in Eastern Europe.

As other nets arrived

As other networks began to arrive (such as EARN and NORDUnet), and the research networking community developed, EUnet joined a variety of cooperative efforts, including participating in the initial discussions which led to the establishment of the CERN-Amsterdam link consortium, which included EARN, IBM/EASInet, CERN and NORDUnet. EUnet was also involved in backup and other service agreements with other international networks such as NORDUnet (from whom we obtained our Nordic connectivity for several years) which resulted in what was known as the Stockholm-Amsterdam link. EUnet was present in the initial discussions which led to the formation of the EBONE, which stood in large part on the bandwidth and membership of the CERN-Amsterdam link consortium and the Stockholm-Amsterdam link. EUnet's Daniel Karrenberg (now leading the RIPE NCC, see elsewhere in this issue) made the original proposal to form RIPE, which provides a forum for cooperation between European IP providers.

Early availability and impact

EUnet has been either the first or among the first Internet-related service providers in nearly every country in which it operates. In particular, EUnet is proud to have provided early access to Bulgaria, Hungary, Czechoslovakia, Yugoslavia, and the Soviet Union (so early that most of those countries no longer exist).

During the failed Soviet coup of 1991, EUnet, through its local provider RELCOM, was a key source of public information within the Soviet Union and to the outside world. This was due to the fact that many western news agencies used EUnet to transfer information between the Soviet Union and the rest of the world.

The commercialization of EUnet

By 1990, EUnet recognized the need for major change. The growing success of the Internet in general and EUnet in particular meant that the EUnet NOCs were outgrowing their home institutes, and that a new organizational approach to running the network was required. The growing requirement for equipment and leased lines resulted in a level of risk and need for capital which could not be supported under academic or user group umbrellas. A European-level manager was hired in the beginning of 1991, and analysis and planning were undertaken to determine the shape of EUnet in the 1990s.

Profile: EUnet (continued)

In 1992 it was agreed that EUnet would be formally constituted as a commercial company, which has now been accomplished at two levels: The national EUnet service providers moved from their institutional homes and reorganized themselves as commercial companies and, in December 1992, EUnet Limited of Ireland, a for-profit company, was formed. All European-level activities within EUnet are performed under EUnet Limited, which is owned by the national EUnets and EurOpen (formerly EUUG) which holds a minority position.

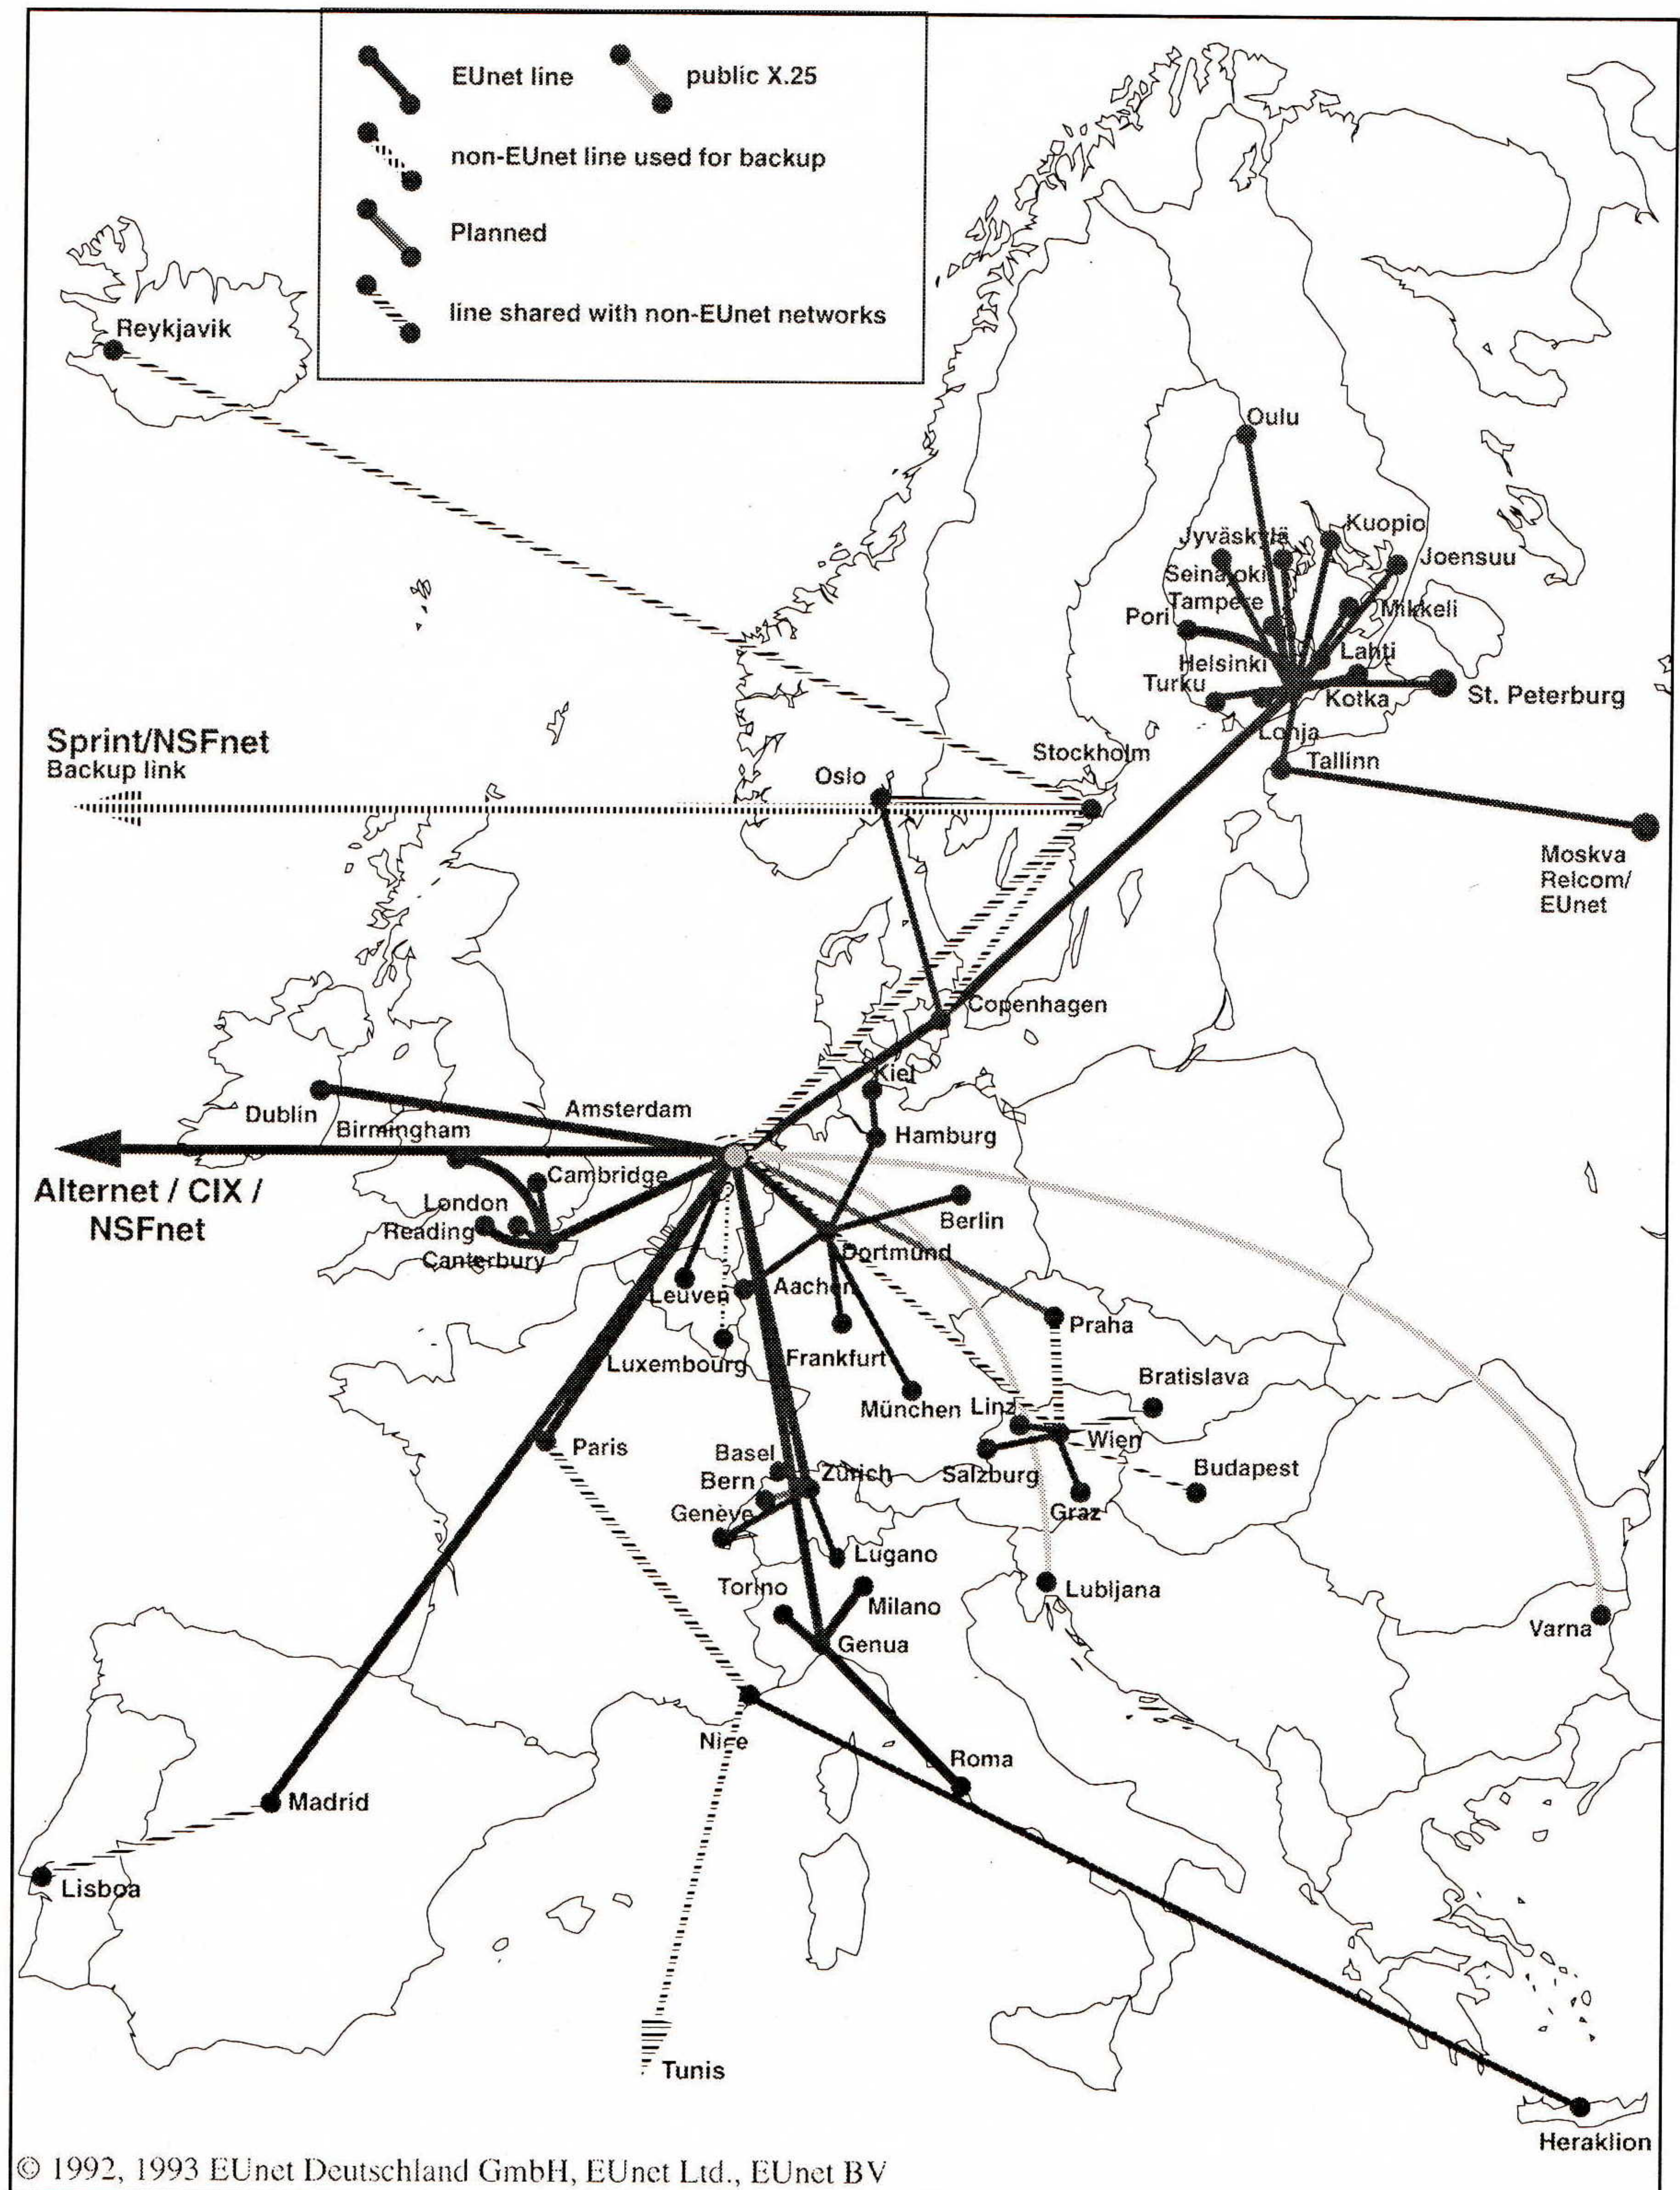


Figure 1: European EUnet infrastructure

Business management structure

The earlier informal decision-making process has been replaced by a more regular business structure. There is a Board of Directors, elected by the shareholders. An informal, thrice-annual "backbone meeting" has been replaced by separate technical and business meetings. The technical meeting is attended by European and national level technical leaders from the national EUnets, who address issues such as architecture, operations, security, and new service development. As required, working groups are spun off to execute well-defined, fixed-duration tasks. The business meeting is attended by European and national EUnet managers and marketing staff, who address issues such as internal cost recovery, strategy, and business relations with suppliers and other providers. Day-to-day business management is performed by a chief executive with the support of an executive committee and staff.

Infrastructure independence

In planning for the commercialization of EUnet, three alternatives were investigated for our long-haul traffic requirements: using the proposed *Operational Unit* services (now using the name *EuropaNet*), the EBONE, or to significantly expand upon our own router network. Use of EuropaNet was rejected when we discovered that the funding organization for EuropaNet would not allow any but the government-sponsored national research networks to take an equity position in that network. Use of EBONE was rejected for two reasons: first, EBONE's uncertain status made it inappropriate for use as the basis of a commercial service; second, it became apparent during 1992 that the EBONE, being largely made up of government-sponsored research networks, could not possibly share the interests of EUnet in the long run, and hence could not be expected to respond effectively to the requirements of our growing non-research subscriber base. As a result of this analysis, EUnet decided in 1992 to complete the process of developing its own infrastructure.

EUnet infrastructure now extends across nearly the entire EUnet system, including a fractional T1 from the Amsterdam NOC to UUnet in the United States (and on from there to the *Commercial Internet eXchange* (CIX), and the NSFNET), the cost of which is partially supported by UUnet Technologies. EUnet maintains its own internal leased lines from Amsterdam to nearly all European EUnets, and this summer installed the first commercial Internet terrestrial link to Moscow (via EUnet NOCs in Copenhagen, Helsinki, and Tallinn). The majority of EUnet internal links are today either 64 or 128Kbps. Traffic and bandwidth on most edges is doubling roughly every 6–8 months. The EUnet core, consisting of lines and router equipment at the EUnet NOCs, is made up almost exclusively of Cisco routers primarily running the BGP3 routing protocol.

EUnet's decision to develop its own infrastructure enables us to obtain an appropriate degree of control to ensure the level of service quality demanded by our subscribers.

Wholesale-Retail structure

EUnet today envisions itself as a commercial wholesale-retail network. That is, EUnet provides three vital components of Internet-related service provision for its subscribers: local customer service, extensive infrastructure, and managed access to the world-wide Internet. EUnet has fully-manned network service bureaus in each country in which it operates, and at the European level in Amsterdam. Each of our service teams is "home grown" and has years of experience; they are familiar with their national networking environment including regulations, the other providers, PTTs, and suppliers.

Profile: EUnet (continued)

EUnet operates its own infrastructure both intra-EUnet and to other key service providers and interconnection points. Finally, EUnet manages internetworking relationships with the other players in the Internet and in the network services marketplace through a variety of bilateral and group agreements.

Availability

EUnet services are available to any organizations or entities within the EUnet operating region. We view any restrictions on use and users besides those regarding "good network citizenship" (for example, correct technical operation of equipment) or within the scope of local law, to have done real harm to the development of the Internet. Any constraint which reduces the potential customer base of the Internet holds back the involvement of communities which could greatly benefit from networking, as well as to keep providers more dependent upon non-market sources of income.

**Funding, Economics
and Cost Recovery**

Each EUnet layer must pay for itself (national and international), each service must pay for itself, and each subscribing user (or user organization) must pay for the services they use. At first blush, this may appear so obvious as to not require stating, but in the European Internet it is not entirely common.

EUnet has generally avoided the European networking obsession with scale-before-demand and government funding. Most government-sponsored networks seek the holy grail of high-bandwidth "pipes" (from 1990 to 1992 this mythically-important bandwidth was 2Mbps, it has recently been 34Mbps) and for the government funding required to support it.

The EUnet view has instead been to build infrastructure by accretion: to find ever-increasing numbers of subscribers, each of which will pay for the infrastructure that they actually use. Factors such as EUnet's progress along the learning curve, the slowly-declining costs of telecom bandwidth in Europe, and the economies of scale inherent in networking, allow EUnet to increase bandwidth capabilities ahead of the aggregate requirements of our subscriber base. With our transition to commercial status, this process has been significantly accelerated due to our greater ability to obtain and use capital.

EUnet has generally not pursued significant grant-hunting in the past: major grants can be a dangerous diversion from the business of networking and, if awarded, can have dangerous side effects; they tend to make providers lose track of the critical linkage between internal costs and the prices actually charged to subscribers, and tempt the provider to drift away from a lean internal cost structure. This is reinforced by the typical experience of grant recipients turning into report generators. Internally, EUnet has a cost replacement structure which provides modest temporary discounts to starting EUnets. This follows from our agenda to extend the Internet to new regions and is distinct from the regressive charging model of the EBONE and many other networking organizations.

**Development of the
periphery**

Besides our core networks, EUnet is currently bringing on board new service teams in Algeria, Poland, Egypt, and Romania. These groups are beginning with dial-up connections and are expected to advance to leased-line connectivity in the near future. When bringing on board new nets such as these, EUnet frequently loans equipment in order to increase the pace of connectivity and service development. We are also happy to advance networking by other means: EUnet today donates free transatlantic e-mail transit to FidoNet.

What catches mice?

The European networking community of the 1980s (and perhaps even today) has been dominated by the industrial policy notions embodied in the push for the ISO OSI networking model. EUnet has fortunately never suffered from this variety of religion. We like to take a position similar to that taken by Chinese Premier Deng Xiao Peng, who revolutionized that economy under the slogan (we paraphrase): "Who cares if the packets are black (OSI?) or white (IP?), as long as they deliver data?" We select technology as a function of subscriber demand, usability, economy, and interoperability. Our choice of IP technology was not so much rare wisdom as it was a consequence of the discipline imposed by our self-funding approach.

A view of the future of Networking Economics

Now that the Internet has grown to be of significant size, it has attracted the attention of professional economists, a few of whom argue that theory shows that the Internet cannot possibly be a market. This is a terribly sad conclusion, again reinforcing Economics' reputation as the dismal science. If the Internet cannot be a market, a wonderful opportunity for wide-open competition and the innovation, quality, availability, economy, and transforming power that it promises will be lost. There is of course no way to predict how this will finally play out, but the Roman Rule comes to mind: "One who argues that something cannot be done should take care to stay out of the way of the one who is doing it." In the long run, the choice for the Internet is simple: the same high-cost, bad-service model enjoyed by monopoly service subscribers, or the chance for a reshaping of economic life around a fundamentally new relationship between distance, observation, and action.

Closing

EUnet demonstrates that an effectively large organization may be constructed by like-minded technologists and business people without the supposedly requisite "big science" approach of government. Modern size- and price-reduced computer and communications equipment which have so dramatically popularized computing, have created new freedoms in exceptionally wide area networking. EUnet and other similarly structured commercial networks are an important vehicle for uncovering the full potential of modern networking technology and the promise it holds for improving relations between people, and to their work.

References

- [1] Goldstein, S. & Michau, C., "Convergence of European and North American Research and Academic Networking," *ConneXions*, Volume 5, No. 4, April 1991.
- [2] Stockman, B., "Current Status on Networking in Europe," *ConneXions*, Volume 5, No. 7, July 1991.
- [3] Stockman, B., "EBONE, The European Internet Backbone," *ConneXions*, Volume 7, No. 5, May 1993.
- [4] *ConneXions*, Volume 7, No. 5, May 1993, "Special Issue: Focus on Europe."
- [5] Guy Almes, Peter Ford, and Peter Lothberg, "Proposal for Global Internet Connectivity," June 1992.

*Learn more:
Session IT-4
Thursday at
10:30 am.*

GLENN KOWACK is Chief Executive of EUnet Limited, and has been with EUnet since early 1991. He holds degrees in Mathematics and Psychology from the University of Illinois at Urbana-Champaign. During the 1980s he was a Senior Director at Gould Electronics' Computer Systems Division, where he directed UNIX operating system and related software development. He has also been a producer and presenter of radio programs and was the Chairman of Prairie Air, Inc., which operates FM station WEFT. He can be reached as: glenn@EU.net


MICE: Multimedia Integrated Communication for Europe

by Jon Crowcroft, UCL

Introduction

Multimedia International Conferencing for European Researchers (MICE) is a project funded by the CEC for the duration of 1993. Its purpose is to develop systems for pilot multimedia conferencing for researchers around the community, using (and inventing) standards as need be.

Of Mice and Men



Partners in MICE and connectivity are:

UCL	London, UK	UK/US Pipe, EBONE, EuropaNet, ISDN
INRIA	Sophia, France	EBONE
GMD	Darmstadt, Germany	EuropaNet
RUS	Stuttgart, Germany	EuropaNet
NTR	Kjeller, Norway	ISDN
Oslo U.	Oslo, Norway	EBONE
SICS	Stockholm, Sweden	EBONE

The mice at UCL are:

Charles Eastal	Project Manager
Angela Sasse	Director
Mark Handley	R&D
Stuart Clayman	R&D
Atanu Ghosh	R&D

Peter Kirstein oversees things, with consulting input from Jon Crowcroft and Steve Wilbur. (All UCL Computer Science Department staff and students can be reached using the generic E-mail address: **I.Lastname**, where **I** = Initial of given name, and **Lastname** is family name or equivalent. For example: **J.Crowcroft@cs.ucl.ac.uk**)

Mighty and Mickey Mice

MICE addresses both *conference rooms* based teleworking and *workstations*, albeit with similar communications technology for both. Conference rooms are designed for meetings to be held with significant numbers of people at each room location. They tend to have higher quality audio/video kit, and presentation devices. Workstations are frequently equipped with low cost mikes and cameras. Presentation tools are whatever runs under whatever window system is used at a workstation, together with tools for sharing (either the window display, or the output (e.g., *PostScript*) or display programs that explicitly replicate their output, or shared input to different display programs). All the workstation technology is equally applicable to conference rooms, except that hardware codecs may be used to get more compression leverage for the higher resolution picture needed when there are more people and analog displays (e.g., "whiteboards") within view.

Getting our Mouse in order

So far, MICE has produced:

- *H.261 software and hardware interworking*: H.261 is a TSS (was CCITT) standard for compressed digital video. It is based around *Discrete Cosine Transform* (DCT), and some simple frame differencing procedures. DCT and video frame differencing are moderately compute intensive operations.

This has been based on INRIA's software codec, *IVS*, and BT/GPT hardware video codecs. Much blood and sweat has been shed stripping and cladding packets with H.221 framing needed to talk to various codecs (including doing CCITT CRCs in software!). A variety of other H.261 codecs have been tested for conformance along the way.

Codecs

A *codec* is a special purpose piece of hardware that contains custom processing to *CO*de and *DE*Code video (and/or audio). In particular, it is usually a device that has simple custom parallel hardware (and multiported memory) to be able to run DCT over blocks in a single video frame and to difference several video frames. Most codecs currently implement H.261, although there are other (better) algorithms (and even standards, e.g., MPEG!) available.

Typically, H.261 is then framed using H.221 (one of the more convoluted protocols devised) possibly multiplexed (entwined) with audio, and then transmitted and received over a synchronous bit stream. This was anticipated as being ISDN ($p \times 64\text{Kbps}$ —where p is not 1, the H.221 stripes the data, to provide minimal possible delay for any audio stream through the overall data stream in the usual paranoid telephony engineers way). Of course, the Internet service as provided by IP (whether from NSFNET, PSI or any other provider) is not a synchronous bit stream. Thus we have two tasks to achieve when interworking between hardware and software H.261: 1) stripping and re-encapsulating H.221. 2) providing a “simulated” bitstream environment.

Task 1) is mainly hard work. Task 2) requires smarts: TCP, for instance, provides a bitstream (well, *bytestream*, but what’s 8 between friendly mice?). However, TCP does this by introducing variations in delay through retransmissions and so forth. In fact, by providing a sufficiently large receive window, we did make this work with TCP (20 seconds worth of receive to account for the maximum possible number of retransmits and backoffs between London and Bonn in one demo!). We have to use this as the basis of a playout buffer. Unlike Van Jacobson’s audio tool, *Vat*, we cannot adapt the playout buffer since the video doesn’t have silence etc. Another scheme was introduced so that we could use UDP and UDP on IP multicast: We simply track the video frame boundaries, at the receiver and adjust to loss (rather than retransmit) by padding with previous frame data. This is hard work (and still requires a small receive buffer). It is now working and means that hardware codecs can be transmitters and/or receivers in IVS Video-casts just like IVS software ones.

- *Packets and Circuits*: Being Europeans, MICE makes use of the quite widely available (basic rate) ISDN that is becoming quite popular. This can be used for direct hardware codec interconnection at 64 or 128Kbps, or for IP router connectivity to provide guaranteed bandwidth while the rest of the European Internet infrastructure is still without this functionality. [We are working on fixing that, though]. Meanwhile, for shared applications, UNIX *talk*, audio and standard applications, the continually improving European infrastructure for IP suffices.

- *CMMC*: Working with packets, circuits, soft video and hardware codecs, with a variety of sources and sinks, multicast and unicast, has led to the development of a large system for managing this. The *Conference Management and Multiplexing Centre* (CMMC) is such a system. Basically, when presented with a plethora of requirements for multicast/multiplex, software to hardware video conversion and so on, it plans a set of translations. It is very similar to the algorithm at the core of the PP mail system that works out what types of body part replacements will be needed when going from Multimedia inbound message of one set of standards and protocols to another (akin to, but not quite, Warshall’s transitive closure calculation).

MICE (continued)

- *Specifications*: A key output of the project is the set of specifications for workstation and conferencing room facilities. These will be made available later as they mature. They make an ideal shopping list for future multimedia rodents...
- *Shared Workspaces*: One of the partners is working on shared workspaces—indeed it is a common finding of that audio plus whiteboard is the key combination to get things done, and that video is really a cherry on the icing to impress the managers—when the cats are away, the cameras get a rest.

Scuttling around the Internet Wainscot

An important part of the effort, with a lot of help from the RIPE crew, is European MBONE engineering. A key lesson from all multimedia work is that we need networks that provide more predictable behaviour than the current Internet.

ITU folks refer to this as *Quality of Service*, but in fact it is not a “peculiar and essential character” (see *Webster*), more a simple quantitative measure of a few parameters: *throughput*, *delay* and *loss*. Since we see multimedia as being the predominant traffic source in the very near future (if not today!), and that there only appears to be a small number of services needed, we prefer to think of them as different services with numerical requirements—they all require the same quality of those quantified requirements in the current community. (If we paid by usage, perhaps that might be reflected in a qualifier, such as reliability of quantified services).

For the technology of choice for video and audio, we have two fixed bandwidth requirements that are really pretty straightforward, and a bound for delay and loss. If these are provided or not, that will suffice. The rest of the community should simply see a network with that bandwidth subtracted.

The future

One view of multimedia is that it is the next Big Thing. This “band-wagoning” is understandable in that, like WIMPs, multimedia progress is very *visible* (and audible).

Another view is that it is questionable whether people really want to see their co-workers (or be seen), and that the most that is needed is a phone and a fax machine (or a PC with a built in phone and fax to e-mail conversion facility, preferable all cordless).



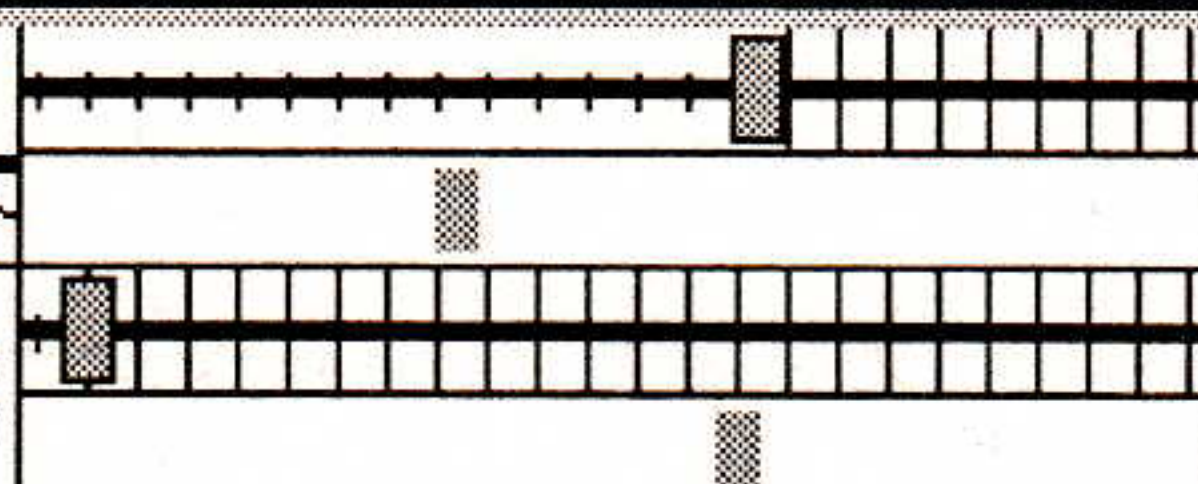
I refer to this latter view as the “deep fried camembert” attitude—that multimedia is just quiche eating and nouvelle cuisine. We believe that in fact, through games, multimedia interactive books, visualisation, distance education and so forth, the requirement for *processable* multimedia will make itself felt as audio and video and unforeseen other media (touch/scent/6th?) become first class objects in the programming world.

When multimedia is truly integrated, collaborative computing will seem as natural as using the TV, the phone and the workprocessor. Otherwise, we may as well go back to lighting fires on hilltops and quill, ink and parchment.

*Learn more:
Sessions AP-3,
AP-7 and AP-8.
See program
for details.*

The Next page is screen dump from a typical MICE “session” —>

JON CROWCROFT is a Senior Lecturer at University College London where he has been engaged in Internet related research for about 10 years. He got his BA from Cambridge some time ago, his Masters from London more recently, and intends completing his PhD before Cyberspace is fully colonized. He can be reached as: J.Crowcroft@cs.ucl.ac.uk

		xterm	
	<pre> Password: 1 danger,jenc93,uninet.no# jove /etc/mrouted.conf 2 danger,jenc93,uninet.no# ps auxw fgrep mro root 105 0.0 0.9 68 272 ? S 07:12 0:00 /usr/multicast/mrouted.n ew root 373 0.0 0.7 132 192 p2 S 08:54 0:00 fgrep mro 3 danger,jenc93,uninet.no# cd /usr/multicast 4 danger,jenc93,uninet.no# kill 105;./mrouted.new 5 danger,jenc93,uninet.no# ./minfo danger 129,241,160,18 (danger,jenc93,uninet.no) [version 1]: 129,241,160,18 -> 129,241,185,10 {f1-1,jenc93,uninet.no} [1/1/tunnel] 129,241,160,18 -> 129,240,150,150 {white,uto.no} [1/64/tunnel/srcrt] 6 danger,jenc93,uninet.no# ./minfo Received a unknown RemoveHost Packet Received a unknown RemoveHost Packet Usage: minfo [-t timeout] [-r retries] router 7 danger,jenc93,uninet.no# screendump ~mhandley/dump 8 danger,jenc93,uninet.no# xli ~mhandley/dump /usr/home/mhandley/dump is a standard 1152x900 8 plane color Sun rasterfile Compressing colormap...129 unique colors Building XImage....done 9 danger,jenc93,uninet.no# screendump ~mhandley/dump 10 danger,jenc93,uninet.no# screendump ~mhandley/dump </pre>		
	<p>224.5.17.11/3456</p> <p>224.5.17.11/3456</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Mark Handley@JENC, Trondheim, Norway <input checked="" type="checkbox"/> Van Jacobson (LBL) <input type="checkbox"/> Claus-Dieter , RUS Stuttgart <input type="checkbox"/> Atanu@uci <input type="checkbox"/> for the luv of mike:-) <input checked="" type="checkbox"/> Hans Eriksson, SIC <input checked="" type="checkbox"/> prunes? <input type="checkbox"/> Blaise @ Inria 	 	
		LBL Visual Audio Tool v2.8alpha	

mic@f1-1.jenc93.uninett.no Super CIF

Conclusions

- Congestion collapse can be avoided
- The proposed channel service algorithm in conjunction with X.25 built-in flow control mechanisms, can avoid the congestion collapse in RC relays, even under severe load situations
- Fairness can be achieved

total transit delay QoS parameters are respected even when the above the system bottleneck capacity

The proposed algorithm creates "firewall"...

Running OSI applications over a CLNS network

by Colin Robbins, NeXor and Paul Barker, UCL

Introduction

In the May 1993 edition of *ConneXions*, we wrote an article entitled "You cannot promote OSI applications over OSI networks." In this we recalled frustrating experiences we'd endured whilst attempting to demonstrate OSI applications over X.25 and CONS technology at various conferences. We concluded that X.25 was often badly implemented and hard to configure; from the responses received (some of these were very entertaining—thank you!), it seems that our experiences are not unusual.

In the May article, we compared the complexities of configuring a host to use X.25 with the relative ease of configuring the same host to use TCP/IP. We were not surprised at the degree of difference as the TCP/IP networking code is a fundamental part of the operation system being used, whereas the OSI X.25 and CONS implementations were provided as separate packages.

Since May we have had the opportunity to assess whether our criticisms of CONS implementations are also applicable to CLNS implementations. As with the CONS implementations we have experienced, the CLNS implementations used are supplied as an add-on package to the base operating system. Colin, who had to port some OSI applications to run over CLNS, had no previous practical experience of setting up a CLNS network.

Setting up a CLNS LAN running OSI applications

The porting work involved three different manufacturers' machines. Two used the "standard" XTI interface to CLNS, and the third had a proprietary interface. The situation was similar for X.25 and CONS.

There were no cabling problems as the Ethernet cables were already joined for TCP/IP networking.

Configuring the CLNS sub-systems on the machine was relatively easy on each machine. Essentially there was a table or configuration tool in which the local NSAP had to be registered, together with NSAP-to-SNPA (MAC address) information for the other two machines (ES-ES routing was used). That was it. Compare the ease of this with the problems we had with CONS/X.25 in trying to configure logical channels, etc.

When the applications were tested, only minor XTI problems were discovered; these were caused by differences in the "standard" XTI interface between the machines. First, the format of the address buffer used by XTI is not defined, so each machine uses its own format. This problem also affects CONS implementations. Second, one of the XTI interfaces did not support the "automatic address conversion" option. These differences mean various implementations of a "standard" API needed different driving routines. The allowing of options within standards is a pervasive problem, and invariably leads to porting and interworking problems.

Better supported

The conclusion we draw, based on an admittedly small sample, is that CLNS implementations are better implemented. Picking up the theme of our earlier paper, we suspect that this is probably because CLNS is taken more seriously than CONS by the manufacturers as they have more pressure from customers to provide good quality software.

The manufacturers technical support is good too. We had a problem with one of the XTI interfaces returning an incorrect error code during asynchronous connection establishment. The problem was quickly identified and remedied by the manufactures engineers.

CONS and CLNS interworking

There are still routing problems to be solved before CLNS can be used in a WAN on the Internet scale, but we believe that solutions are inevitable so long as there is customer and manufacturer interest.

However, one serious problem remains, which stems from there being two OSI network protocols. Following the CLNS porting work described above, the applications can now be run with a dual OSI network stack—both CONS and CLNS. However, both stacks use the same NSAP addressing scheme. This poses a problem: when an application passes an NSAP address to the transport layer, how does the transport layer know which network stack to use?

It cannot tell from the NSAP, since both CONS and CLNS use the same address space. The only way is to use additional information. Local address mapping tables suffice for small networks, but cannot solve the problem in a large OSI environment. Some OSI proponents have suggested that there is no problem as the OSI X.500 Directory can be used to indicate the network type corresponding to a particular NSAP. Whilst the directory could hold this information, it does not solve the problem, but merely defers it.

The X.500 directory is in a similar position to the other applications; when it needs to make a connection to a remote DSA, how does it know whether the NSAP address is a CONS NSAP or CLNS NSAP? In fact, it cannot tell. It cannot look up these addresses in a remote part of the directory, as the remote DSA cannot be reliably contacted until the CONS/CLNS answer is known—a “Catch-22” situation. Directory interworking is seriously compromised as the referrals which are passed between DUAs and DSAs cannot contain any additional network type information without an extension to the X.500 protocol.

It is frustrating that this problem has been discussed by the standards committees, but that no solution has been found. Some of the difficulties arise because it is seen as someone else's problem; the lower layer groups say it is an application problem, whilst the application groups say it is a lower layer issue. An impasse has been reached.

There was a proposed fix for this problem in the 1992 round of standardisation, but the solution was rejected. Irrespective of the reason for the proposal's rejection, the failure to find a solution suggests the OSI committees are more interested in generating a perfect model, rather than making something work.

Solutions

If OSI is to be made to work in a global context, then CONS and CLNS are going to co-exist. Potential solutions include:

- A CONS/CLNS gateway protocol needs to be defined and made to work. Either network must accept *any* NSAP and route it accordingly.
- An NSAP could contain information about the network type used. Application relays could then route between CONS only and CLNS only systems.
- The X.500 protocols “referral” mechanism could be extended to include information about the NSAP addresses passed, allowing X.500 to be used by other applications for address look-up.
- All machines could run dual stack!
- If you have a protocol like X.500 that assumes all servers are connected on the same network, then define only one network service—not two!

OSI applications over a CLNS network (*continued*)

Conclusion

The OSI protocol suite is a high functionality suite, but has attendant complexity. These two articles have demonstrated practical problems in trying to deploy OSI applications on a large scale. These are caused in part by this complexity, but also, in the case of CONS, by poor implementations. We have also noted a failing by standards bodies to solve practical interworking problems which must be addressed if OSI networking is to be widely adopted.

References

- [1] Robbins, C., Barker, P., "You cannot promote OSI applications over OSI networks," *ConneXions*, Volume 7, No. 5, May 1993.
- [2] "Information processing systems—Telecommunications and Information Exchange between systems—Protocol for Providing the Connectionless-mode Network Service," ISO 8473, March 1987.
- [3] Marshall T. Rose and Dwight E. Cass, "ISO Transport Services on top of the TCP," RFC 1006, May 1987.
- [4] Hagens, Rob, "Components of OSI: CLNP or A Day in the life of Ivan CLNPacket," *ConneXions*, Volume 3, No. 10, October 1989.
- [5] Callon, Ross, "An Overview of OSI NSAP Addressing in the Internet," *ConneXions*, Volume 5, No. 12, December 1991.
- [6] Benford, Steve, "Components of OSI: X.500 Directory Services," *ConneXions*, Volume 3, No. 6, June 1989.
- [7] Jacobsen, Ole, "The Trouble with OSI," *ConneXions*, Volume 6, No. 5, May 1992, p 62.
- [8] desJardins, Richard, "*Opinion*: OSI is (Still) a Good Idea," *ConneXions*, Volume 6, No. 6, June 1992, p 33.
- [9] Metcalfe, Bob, Smart, Bob, and Blackshaw, Bob, "Letters to the Editor," *ConneXions*, Volume 6, No. 6, June 1992, p 37.
- [10] Rose, Marshall, "Comments on: '*Opinion*: OSI is (Still) a Good Idea,'" *ConneXions*, Volume 6, No. 8, August 1992, p 20.
- [11] desJardins, Richard, "Comments on: '*Comments on: Opinion*: OSI is (Still) a Good Idea,'" *ConneXions*, Volume 6, No. 10, October 1992, p 43.
- [12] desJardins, Richard, "Internet 2000," *ConneXions*, Volume 6, No. 10, October 1992, p 24.
- [13] Hoffmann, Harald, "Letter to the Editor," *ConneXions*, Volume 6, No. 11, November 1992, p 31.
- [14] Crocker, David, "Letter to the Editor," *ConneXions*, Volume 7, No. 2, February 1993.

COLIN ROBBINS received a BSc in Computer Science and Electronic Engineering from University College London. Following this he spent 3 years at UCL, where he was the primary implementor of the QUIPU X.500 system, which is widely used in both US and Europe. He is now responsible for the development of Directory and related OSI products at NeXor Ltd. (formerly X-Tel Services Ltd.) and is also active as a consultant with several European Commission funded OSI research and development projects. He can be reached as: c.robbs@nexor.co.uk.

PAUL BARKER received a B.A. in Economics from Exeter University and a Diploma in Computer Science from Birkbeck College, London. He has worked in the department of Computer Science at University College in computing since 1986. During that period he has worked on a variety of projects concerned with X.500 Directory Services, and is currently investigating the use of the Directory for bibliographic purposes. He is past secretary to the U.K. Academic Directory Group. He can be reached as: p.barker@cs.ucl.ac.uk.

*Learn more:
Tutorial T21
Monday and
Tuesday, and
of course "The
Great Debate:
OSI vs. TCP"
on Thursday*

Is DFS the Distributed File System of the 90s ?

by Brice MuangKhot, IBM France

Abstract

Nowadays, NFS (*Network File System*) is data sharing's *de facto* standard. Yet NFS's authentic challenger is DCE/DFS (*Distributed Computing Environment/Distributed File System*). DCE/DFS is based on Transarc Corporation's AFS (*Andrew File System*), and hence, inherits all of its functionality. It is thus similar to AFS in a number of ways. However, DFS, unlike AFS, is no longer a stand-alone product. It is fully integrated with all of DCE's components. DCE/DFS could emerge as the distributed file system of the 90s. This article provides an overview of the power of an emerging distributed file system technology. It also tries a comparison with the NFS *de facto* standard.

Introduction

File sharing between users on a network was one of the first distributed computing applications. The need to share data was born with the fantastic LAN boom in the beginning of the 1980s. Sun's NFS is undoubtedly one of those data sharing applications, as, of course, is AT&T's RFS (*Remote File System*). NFS has been shipped by Sun Microsystems since 1984. The distributed file technology that is at the basis of the DFS component of the Open Software Foundation's DCE (*Distributed Computing Environment*) is the well-known AFS (*Andrew File System*). AFS was co-developed by Carnegie Mellon University (CMU) and IBM in 1983. It provides the same file sharing facilities as NFS but with several enhancements. AFS's main qualities are its *scalability* (it is designed to adapt itself to a network's growth up to over 5000 workstations), *security* (data is protected from non authorized users), *availability* (data can be replicated on many machines at the same time), *performance* (data is cached locally on the client machine before it can be accessed), *data integrity* (synchronization of data update is done using the token mechanism). All this is unlikely to be accidental, nor could these functions have been added afterwards. They all had to be taken into account at the beginning of the design phase. DCE/DFS is based on AFS version 4. AFS was originally developed at CMU, and evolved into a product as the result of work done by Transarc Corporation.

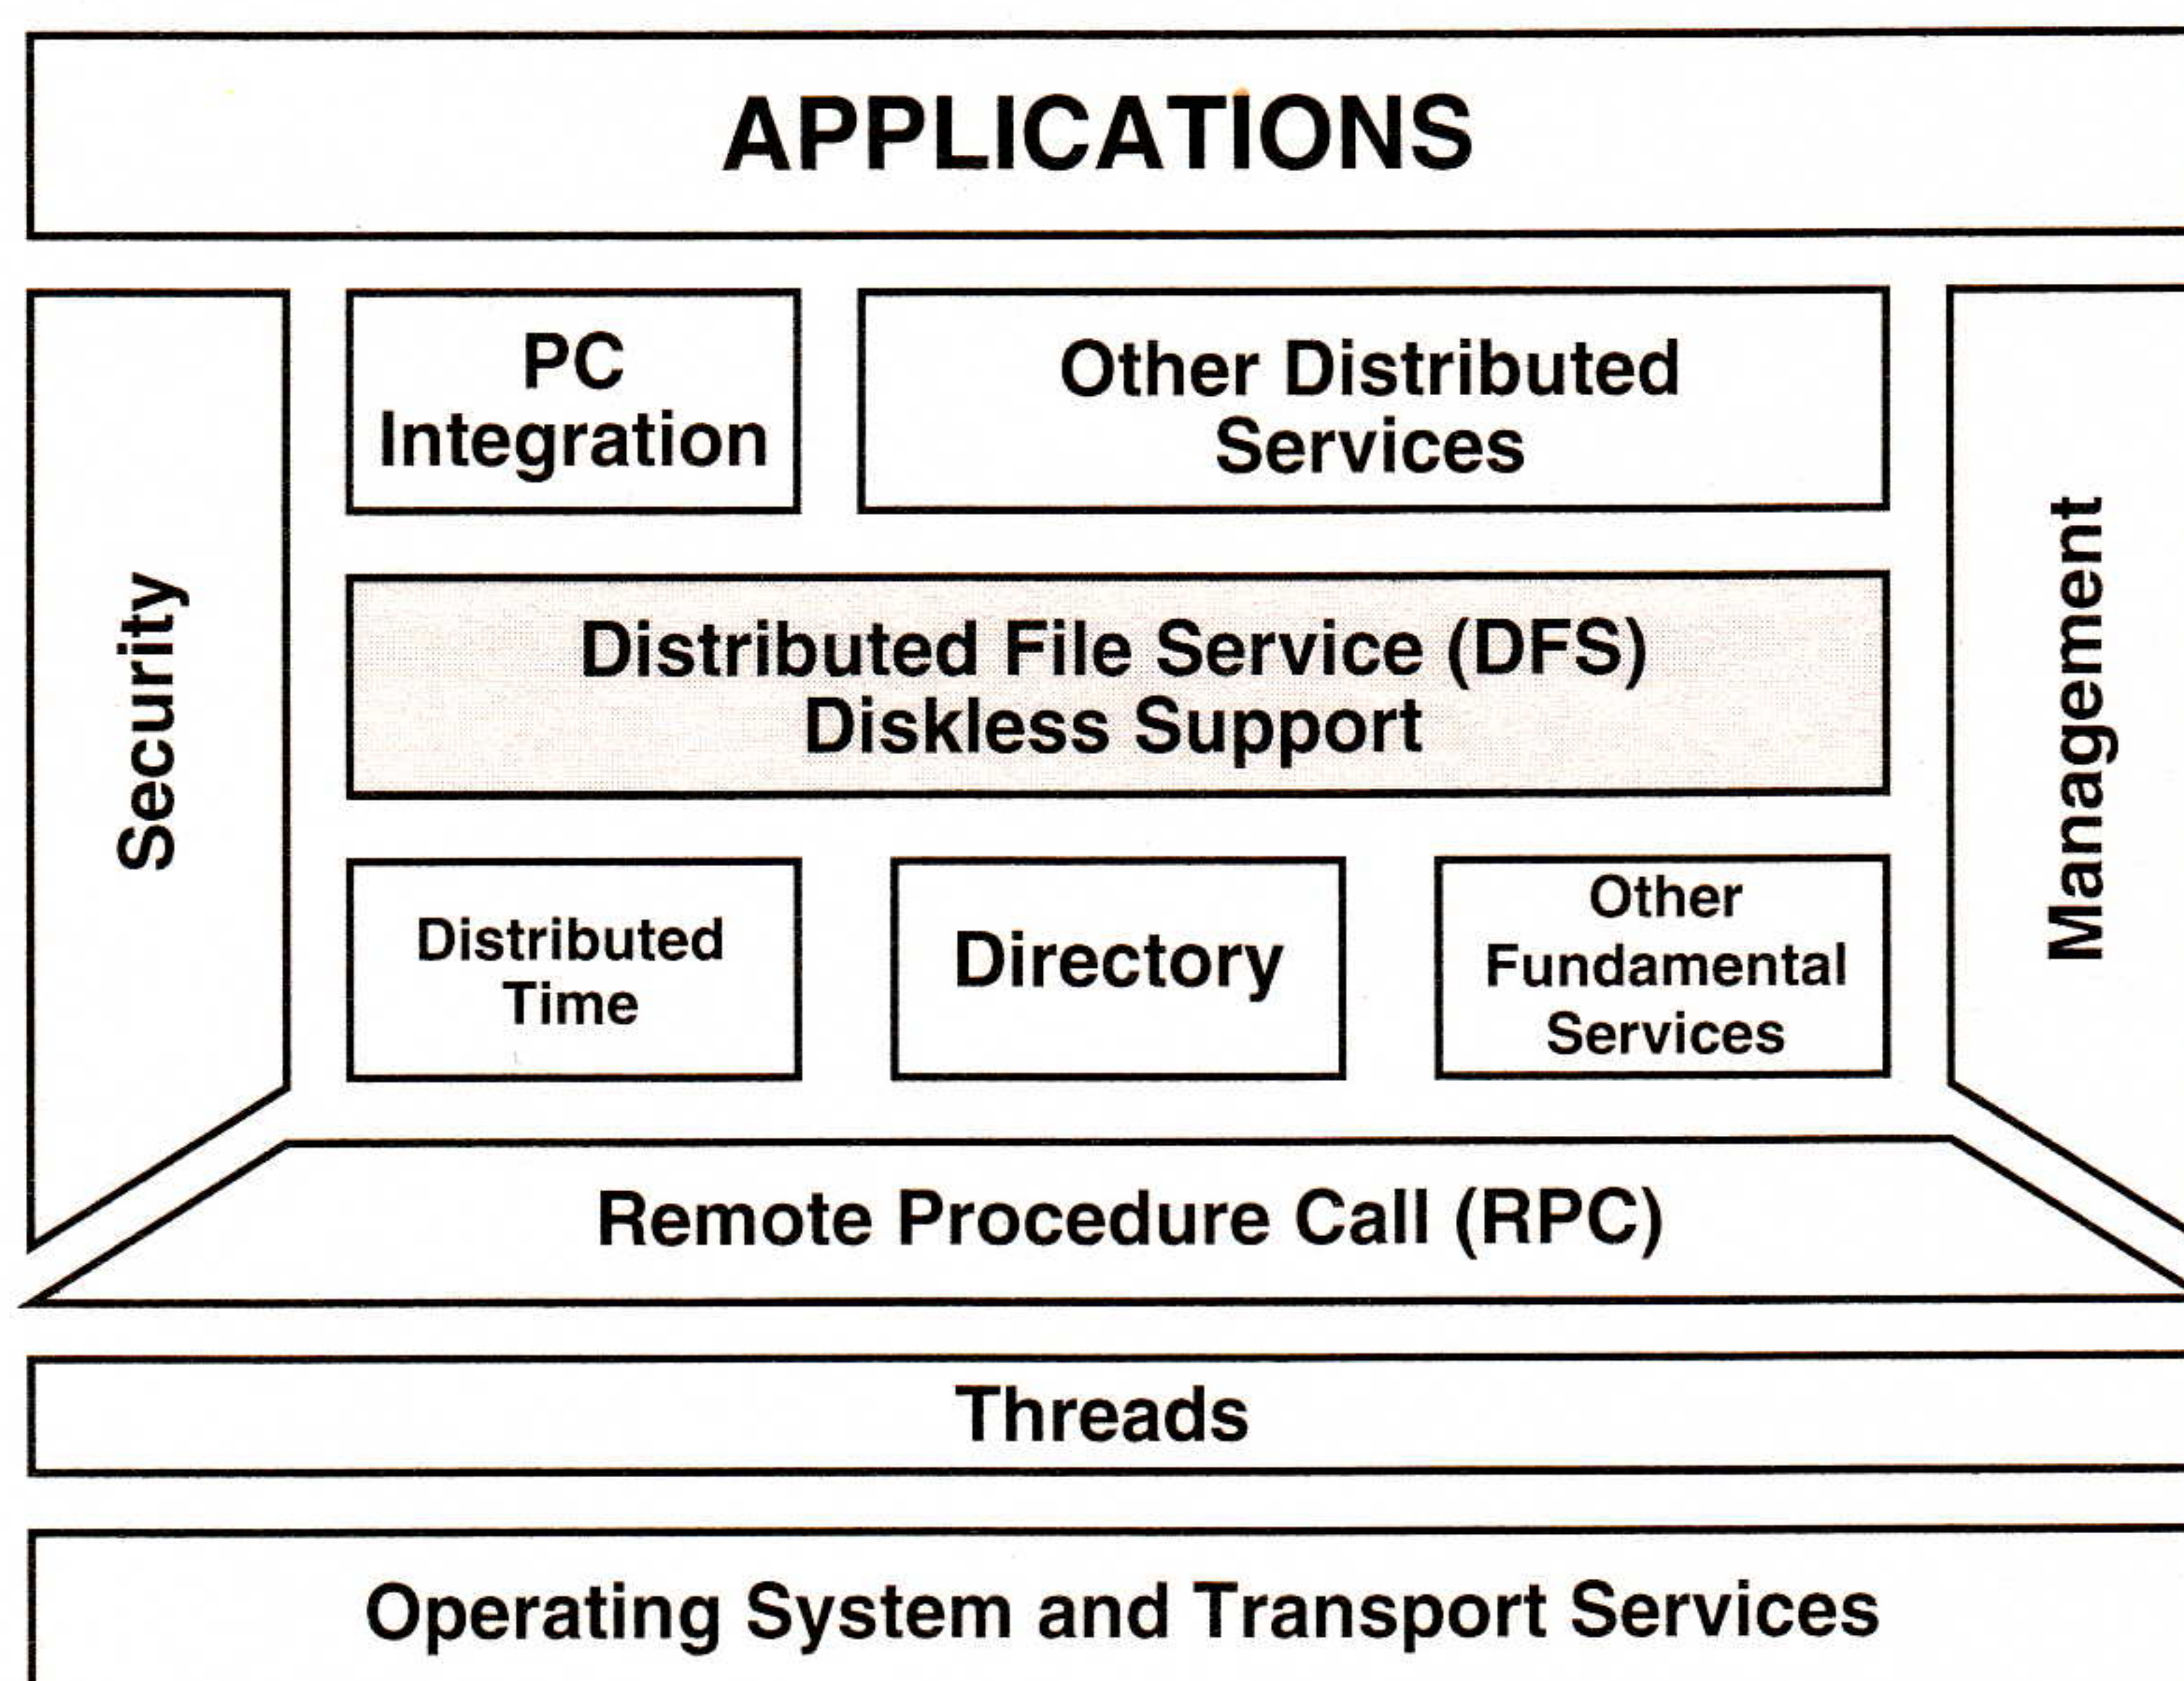


Figure 1: OSF/DCE Architecture

DFS (continued)

What is DFS?

DFS provides data sharing services for use in a distributed environment by extending the local file system model to remote systems. NFS provides the ability to store and access data at remote locations, typically via a client-server model. DCE/DFS is an example of an application that takes advantage of all the features of a state of the art distributed computing environment like OSF/DCE. DCE and DFS, integrated together, provide high availability of data and resources. DCE/DFS inherits all of AFS's attributes but, contrary to AFS, it is fully integrated with the other DCE components. (See Figure 1).

We now take a closer look at DFS's different features and advantages.

Availability

DFS achieves high availability through the *replication mechanism*: there may be more than one file server hosting a given file, so that if one server is down, the same file can be accessed through another server. What's more, copies of files are locally cached before they can be accessed, so that even if the client is temporarily disconnected from the network, it may still be possible for a user to get to a copy of a file in the local cache. In addition, administration tasks do not make file services unavailable. DFS administration can take place as users continue accessing files. Both backup and relocation of DFS files can be done on open files. This feature, again, increases high availability.

Uniform file access and transparency

DFS is based on the DCE *Global Name Space*. A DFS file can be accessed using a unique name, no matter where in the distributed system it is accessed from. Users need not know the network address or name of the file server(s) where the file is hosted. Access to a file is thus completely transparent.

Transport independency

DFS is based on the DCE base services. DCE is designed to be independent from transport and network layers (TCP/UDP, TP/OSI, IP, X25, DECnet, etc.). Furthermore, DCE is designed to work not only over a LAN, but over a WAN too. A product that can work over wide area networks provides important functionalities and advantages. These abilities are achieved through the use of a *local cache mechanism*: DFS, thus, generates minimum network traffic. Despite its present popularity, NFS, unlike DFS, can only work with the TCP/IP protocol suite (and particularly with UDP). NFS is mainly designed to work with LANs. It is not advisable to use NFS over WANs.

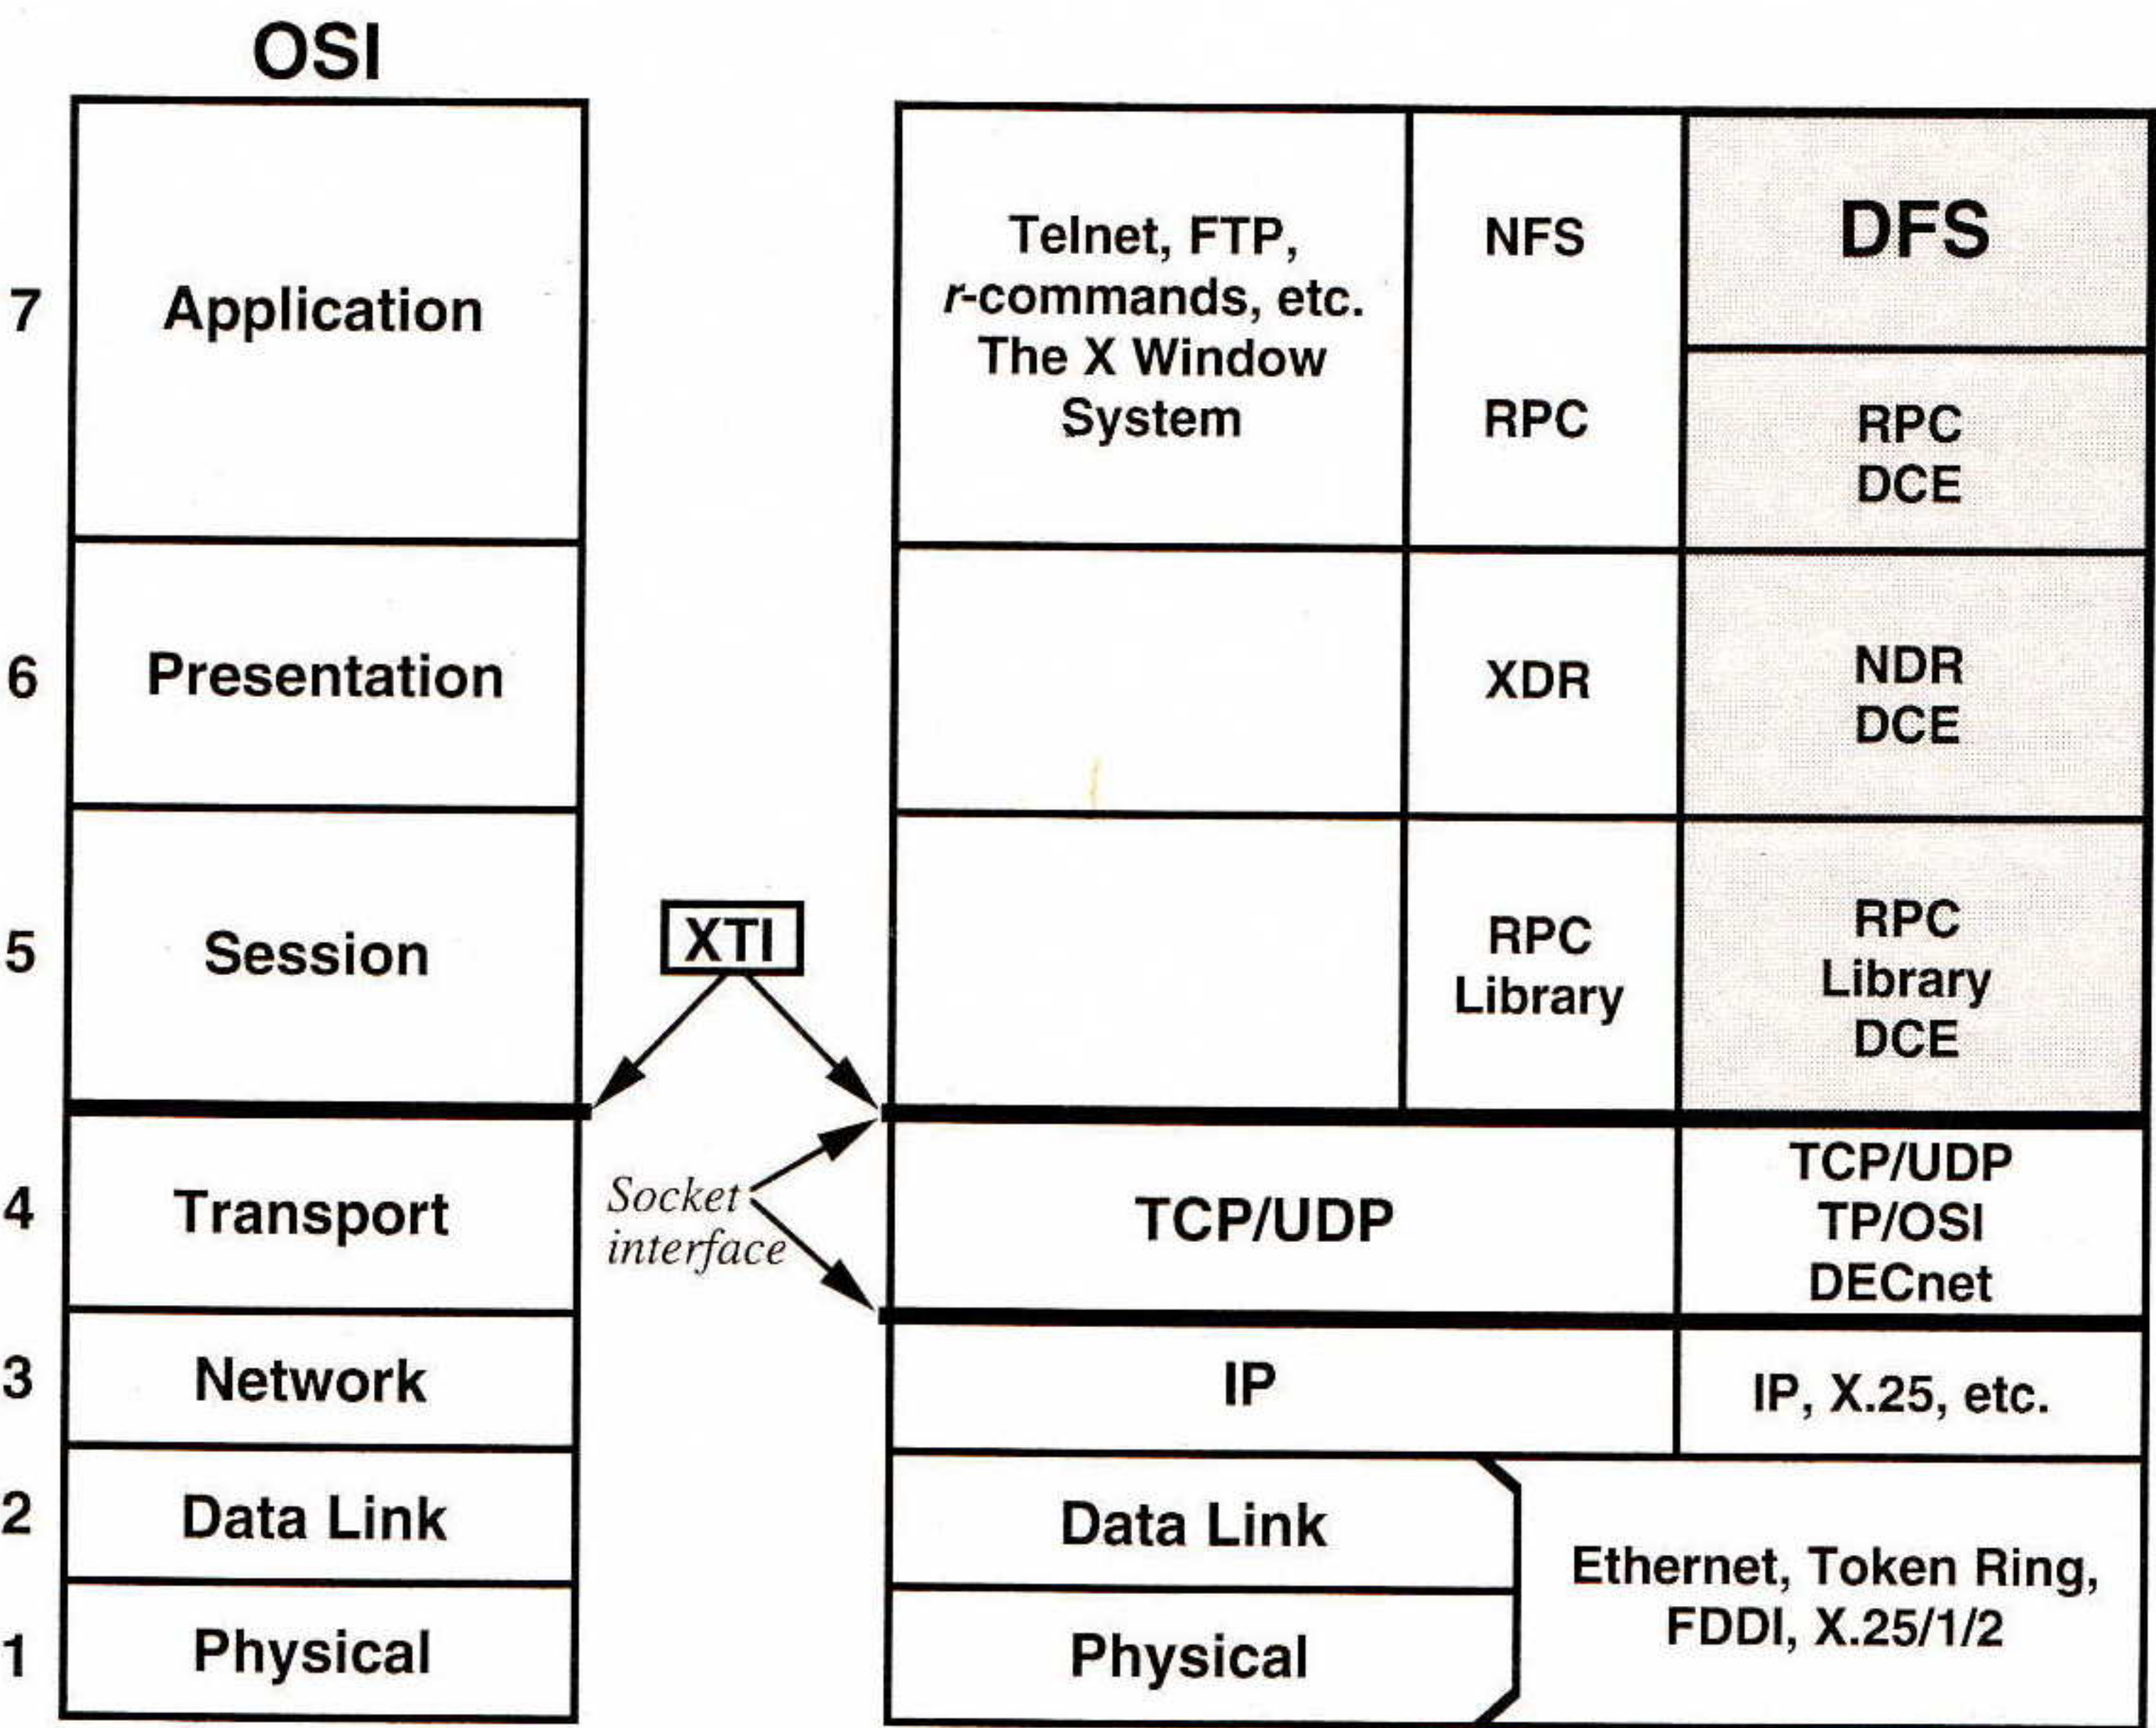


Figure 2: Transport Layer Transparency

Performance

DFS's high performance file service results from the use of the cache mechanism: Fast response is partly achieved through the caching of file and directory data on the client machine. This reduces file access time and network traffic. NFS, on the other hand, accesses the remote server at regular intervals. (3 seconds for files, 30 for directories). Unlike NFS, DFS only accesses the remote server when data is to be updated. This helps reduce unnecessary network traffic.

Data integrity

Unlike NFS's stateless implementation, DFS exploits *statefulness*: the server keeps track of all the exchanges with the client using the *token* mechanism. Hence, the server always knows which clients have cached copies of a particular file. Before accepting an update request, a server revokes the read or write tokens from the clients using the same file. So DFS can ensure that users are always working with the most recent version of a file.

Security

DFS is based on DCE base components: thus, it takes maximum advantage of the security facilities provided by the DCE security service. DCE security provides DFS with user authentication, verification of user privileges, and authorization control (and sometimes, a privacy facility, though there exists restrictions for some countries).

Interoperability with NFS

DFS can interoperate with existing file services like NFS. An NFS client can access files exported by DFS (with some security restrictions). This enables NFS sites to smoothly migrate to DFS.

Standards

DFS is POSIX-compliant. DFS conforms to POSIX 1003.1 for file system semantics, and POSIX 1003.6 for access control security.

Conclusion

NFS is undoubtedly the most popular distributed file system today. But it has many weaknesses that DFS tries to cure: uniform data access, performance, availability, integrated security, data integrity and network transparency. In spite of its being a recent standard, DFS could emerge as *the* distributed file system. DCE and DFS are and will be implemented over various platforms, from PCs to mainframes. DFS has enough potential to become the distributed file system of the 90s, when a user no longer "belongs" to a machine or a network but to a globality, a set of worldwide systems and networks. If Sun's motto is "The Network is the Computer," then, DCE and DFS could enable us to say "The World is the Computer."

*Learn more:
Tutorial T61
Monday and
Tuesday.*

References

- [1] "The Sun Network File System: Design Implementation and Experience," Russel Sandberg, Sun Microsystems.
- [2] "Make or Take: Decision in Andrew," James H. Morris, CMU, USENIX Winter Conference 1988.
- [3] "Scale Performance in a Distributed File System," CMU, *ACM Transactions on Computer Systems*, February 1988.
- [4] "Scalable, Secure, and Highly Available Distributed File Access," Mahadev Satyanarayanan, CMU.
- [5] "File Systems in a Distributed Computing Environment," OSF White paper, September 1990.
- [6] "Introduction to DCE," OSF.
- [7] "DCE/DFS versus AFS," Dawn C. Stokes, IBM Austin.
- [8] Chappell, D. "The OSF Distributed Management Environment," *ConneXions*, Volume 6, No. 10, October 1992.
- [9] Chappell, D. "The OSF Distributed Computing Environment (DCE)," *ConneXions*, Volume 7, No. 3, March 1993.

BRICE MUANGKHOT is an Engineer from Ecole Supérieure d'Informatique (1985), Paris France. He has worked for Bull on a Networking project (FTAM-CTS/WAN) within the ESPRIT Program for the European Community. He has been working for IBM France in the UNIX/AIX Competence Centre since 1989. He has worked on various network projects and is now in charge of second level support of distributed computing products. He is a member of the Networking Group in the Association of French Unix Users (AFUU). He can be reached as: brice@ibm.fr

RENATER: A 34Mbps Network Service

by Christian Michau, RENATER

Introduction

The French national network for research, education and technology, *RENATER*, is now fully operational. After a pilot phase, *RENATER*'s service was started in December 1992 when we established the first implementation of the national backbone and its international gateway giving access to EBONE and the NSFNET.

The extension of the backbone has been conducted very quickly and, according to the initial plans. Phase 1 of the project was completed at the end of June 1993, when we linked all the regions (metropolitan parts of France) to the national backbone. During July 1993, the upgrade of the main North to South axis of *RENATER* to 34Mbps was achieved between Paris, Lyon, Marseille and Montpellier.

The groupement Renater

A "groupement d'intérêt public" (GIP) *RENATER* has been set up by the funding organisations (Ministry of Education, Centre National de la Recherche Scientifique (CNRS), Commissariat à l'Energie Atomique (CEA), CNES, Institut National de Recherche en Informatique et en Automatique (INRIA), and Electricité de France (EDF)) to manage the project, develop associated services and guarantee the quality of services.

The provision of the infrastructure and its management has been contracted to France Telecom. As a first step in the application field, the management of the mail coordination and interconnection is being contracted to the University of Rennes and INRIA. Michel Lartail has been appointed as the director of GIP *RENATER*. The GIP *RENATER* is currently organised into 2 departments:

- The operations international division managed by Christian Michau is in charge of the operations, quality of services and international connectivity,
- The development division, managed by Jacques Prevost, is in charge of the development of the services and the preparation of further evolution.

Architecture

CHRISTIAN MICHAU is a graduate from University Rene Descartes in Paris. Since 1971, he has been working for Centre National de la Recherche Scientifique (CNRS) in the provision of computer and network services. He currently manages the network department of CNRS which is in charge of organizing and developing network services for this main research agency. He has been active in the *RENATER* initiative and is now in charge of the operations group and international services within GIP *RENATER*. In the past 10 years, he has been involved in European network initiatives. He is a member of the EBONE Management Committee.

RENATER is basically an IP service, however a part of the current contract includes the provision of X.25 service by Transpac for specific purposes. Currently, 11 regions in France have established regional networks. For the moment, all of them have been contracted to France Telecom. Each of them has a specific contract for 2, 3 or 4 years, pricing rules, and associated user policy. Some of them may provide services to commercial users as well as the R&D community.

RENATER itself is the backbone interlinking these regional networks and providing the international connectivity. However, end-to-end service for the users has been defined in the contract to ensure the quality of the service provided. In the current terminology, the backbone is named RNI (*Reseau National d'Interconnection*) and the international gateway NTI (*Noeud de Transit International*). The NTI has been very closely integrated to EBONE. In those regions that have not yet implemented regional networks, RNI provides direct connectivity to at least one site, generally an academic site. At the end of July 1993, more than 200 sites were connected to *RENATER*, 47% of them have high speed access points (at 2Mbps or 34Mbps).

The traffic is growing very quickly: Last June, the monthly exchanged traffic between the backbone, the regional networks and the NTI went over 768Gbytes (not including the regional traffic).

Technology

Currently, RENATER is using a classical IP router technology supporting up to 34Mbps circuits. Some of the transit routers are connected on FDDI rings. Generally, when a user is connected to RENATER, France Telecom provides the local router connected through an Ethernet or FDDI to the site router. This is done to preserve management domains of responsibility. Steps are being taken to prepare a smooth migration to an ATM technology within the backbone.

Management

France Telecom is responsible for the management of the topology, the links and the routers. They assume management of the Layer 3 service. This function is provided through a central NOC located at France Telecom offices. Two regional networks have their own NOCs. A hotline service is available. France Telecom reports to the GIP RENATER through regular reports and statistics.

International gateway

RENATER has decided to direct all its international traffic through EBONE. Therefore, it has been actively participating in the development of EBONE, and will continue to do so. Its Paris EBS is currently linked to the London EBS and the Geneva EBS, and it has a 1.5Mbps connection to the NSFNET. Four RBSs use the Paris EBS to access EBONE and a few additional countries are investigating possible connections. The international gateway and the EBS are managed by France Telecom, and the contract includes 24-by-7 coverage all year.

Costs

A site connected to a regional network is charged by France Telecom only for the access to this regional network. Costs associated with the national and international connectivity are paid by the participating organisations of the GIP, or through specific contracts for others.

Cooperative agreement

All the development of RENATER has been established under the umbrella of a cooperative agreement signed by the Ministers Education, Research and Telecommunications in 1991. This provides the context for a close partnership between France Telecom and GIP RENATER. Over the last 9 months, this partnership has proven very successful. The migration from the previous networking facilities to RENATER has gone smoothly, and stable service has been provided to the users.

The future

Issues for the future development of RENATER are outlined below.

- *Quality of Service:* It is part of the current agreement with France Telecom to move to a service oriented contract. Therefore, criteria for quality of service are being defined between GIP RENATER and its contractor and will be implemented within one year. At this moment, the contractor will be responsible for upgrading the network as appropriate to meet the requirements. This will allow for an open infrastructure for additional partners.
- *International relations:* The international connectivity is a key element of the service provided to our community. RENATER has expressed its formal willingness to support actively the development of EBONE, and is fully prepared to develop the current Paris EBS to become an international exchange facility platform.
- *ATM migration:* A firm commitment has been made to prepare for the migration to ATM. Initial technical steps have already been taken to prepare for this evolution and to start pilot activities in this area.

Conclusion

The establishment of RENATER has been a real challenge for all involved. During this implementation phase, we've received a lot of support and help from our friends working all around the world on global interconnectivity. I want to thank them for their cooperation and for the benefits of this cooperation to our user community.

Learn more:
Session IT-7
Thursday at
1:30 pm.

Using APPC and APPN to Integrate LANs and SNA Networks

by Tim Huntley and Susan Schulken, IBM

Introduction

These days, many companies are finding themselves in the networking business, even though that's not what they want to be doing. That's because these companies are trying to interconnect LANs and WANs that were designed and developed separately, and it's not a simple job. As a result, their network administrators overwhelmed with complex, detailed tasks like setting up filters, configuring bridges and routers, updating routing and control tables, defining new users, and removing old users.

In a typical scenario, Company A uses a SNA network to handle business and financial applications. At the same time, they've installed separate LANs for various groups of scientists and engineers. Now, what happens when a scientist on the LAN wants to see budget information on the host system, or users on different LANs need to share project data? The task of enabling these users to access applications and data across networks falls to network administrators. Unfortunately, given the challenges of integration, these administrators spend more time trying to solve daily networking problems than developing plans that address business issues.

Up to now, network administrators have had only a couple of integration approaches to choose from. Compensating for the limitations of either approach means extra work for these administrators. The most popular approach has been to bridge LAN protocols or route them over different types of links, which is complicated and costly. The alternative is to use a traditional SNA protocol throughout the network, yet this approach does not offer the ease-of-use, flexibility, and performance that users demand.

Now, there's another option: *Advanced Program-to-Program Communication* (APPC) and *Advanced Peer-to-Peer Networking* (APPN). Also known as "the new SNA," APPC and APPN were designed to support today's networking and communications needs, including the need to integrate diverse, multiplatform networks. These technologies can help eliminate the complexities and headaches associated with network integration.

Option 1: Routing LAN protocols

LAN-based protocols, such as NetBIOS, IPX, and VINES, were intended for a limited, specific purpose: to enable PCs to communicate with each other over a stand-alone LAN. These protocols were never intended to be used on other computers, such as midrange machines and mainframes, or over other types of networks. So, if you want to extend LAN protocols to other networks in your business, your options are (1) using bridges to connect LANs, or (2) encapsulating and routing the protocols over other types of links.

The LAN-based protocols include several features that are ideally suited for the stand-alone environment, such as regular broadcasts and time-outs. Yet, these same features are often handicaps when you extend the protocols to other environments. Using either option means extra work to compensate for the shortcomings of LAN-based protocols in larger networks. Plus, if you're routing the protocols, you need to consider the drawbacks of encapsulating data.

- *Broadcasts:* Most LAN protocols broadcast information to every node in the network. A node that is sending data may also broadcast to all rings in the network to determine the location of a specific MAC address. These broadcasts don't pose a problem for small LANs.

But, if you start integrating networks, these broadcasts can result in lots of unnecessary network control traffic that can quickly consume bandwidth.

- *Time-outs:* LAN protocols use time-outs to determine whether data needs to be transmitted again or whether an application's partner has moved. Once you start interconnecting networks, time-outs can cause problems. In general, time-outs are often too short for a WAN environment, especially if you're sending data over slow lines or using a satellite hook-up for file transfer. In these environments, the sender detects the time-out and resends the data over and over again. The network becomes crowded with duplicate packets that the receiver will discard. Because of the overload on the network, bandwidth is wasted and applications may not be able to send data.
- *Configuration:* Finally, configuration in a stand-alone LAN environment is relatively easy and straightforward. When you start integrating LANs, however, you'll find that much more configuration is required. If you use a bridge to connect two LANs, for example, you may want to use a filter to limit access between the two networks and eliminate broadcasts between all nodes in the LANs. Setting up this filter requires extra configuration. Later, if you move computers around, additional configuration is required to update the filter. The result is that, in some companies, one or more network administrators do nothing but update configuration information.
- *Encapsulation:* To route LAN traffic over other networks, data must be encapsulated. While encapsulation is effective, it's not without its limitations. With encapsulation, you have to allow for the overhead of additional header information. You also lose information that could be used to assign priority to certain types of data. In other words, while it is preferable to assign a low priority to a file transfer and a high priority to an interactive session, you lose these distinctions when data is encapsulated.
- *Multiple Networks:* Using and maintaining multiple protocols in your network is complex and expensive. When you use a mix of LAN protocols and other protocols, network management becomes much more complicated. Similarly, it is more difficult to deal with network problems, since you can't rely on the error handling services of one protocol to identify problems in parts of the network that use other protocols. To work with the various protocols, network administrators typically need additional training, which is costly and time-consuming. For these reasons, bridging and routing LAN protocols should be considered a temporary solution, at best.

Option 2: Using traditional SNA

For years, SNA has been synonymous with reliability, security, and high performance. Yet, traditional SNA protocols have not provided the flexibility required for a LAN environment. Just as LAN protocols were not meant to run on large machines and large networks, these WAN protocols were not designed for the needs of PCs and LANs. Using traditional SNA throughout your network requires detailed configuration, more limited choices of network designs, and potential performance problems in the LAN environment.

- *Configuration:* One of the most common concerns among companies that rely on SNA protocols is the configuration requirements. While configuration in a stand-alone LAN environment is fairly simple, SNA networks require much more set-up, by comparison.

Using APPC and APPN (continued)

For example, if you add a new user to an SNA network, you have to coordinate naming to avoid conflicts, update routing tables, and update VTAM and NCP definitions. The configuration associated with moving computers in the network is even more complex, since you have to redo the configuration manually, plus you often need to delete old configuration information in multiple places in the network.

- *Flexibility:* Another criticism of traditional SNA is its lack of flexibility in supporting various network designs. SNA was developed to support hierarchical, host-based networks. This approach does not provide the flexibility to support other network topologies, such as token-ring, mesh, or flat peer-to-peer.
- *Performance:* As noted earlier, SNA protocols have an excellent reputation for providing high performance in a WAN environment. Yet, these protocols have not excelled at maintaining these levels of performance in LANs.

LAN protocols offer exceptional performance, since these protocols require only a small amount of routing header information. In addition, LAN protocols are based on the assumption that LANs are highly reliable networks. This type of assumption can't be made if you're routing information over SDLC or async. Because traditional SNA protocols include additional header information, as well as small send pacing windows, you'll notice a reduction in performance.

Option 3: APPC and APPN

The basic problem with the LAN and WAN approaches is that both were developed for specific environments, and many adjustments are required to make them work outside those environments. APPC and APPN, by contrast, are more generalized protocols; they were designed to function well in small and large networks, over slow and fast links.

APPC supports any-to-any communication between programs on many different types of computers, from the smallest PCs to the largest mainframes. Likewise, APPN was developed to provide intelligent routing of APPC traffic over both LANs and WANs. The combination of these two technologies offers the best of both WAN and LAN computing.

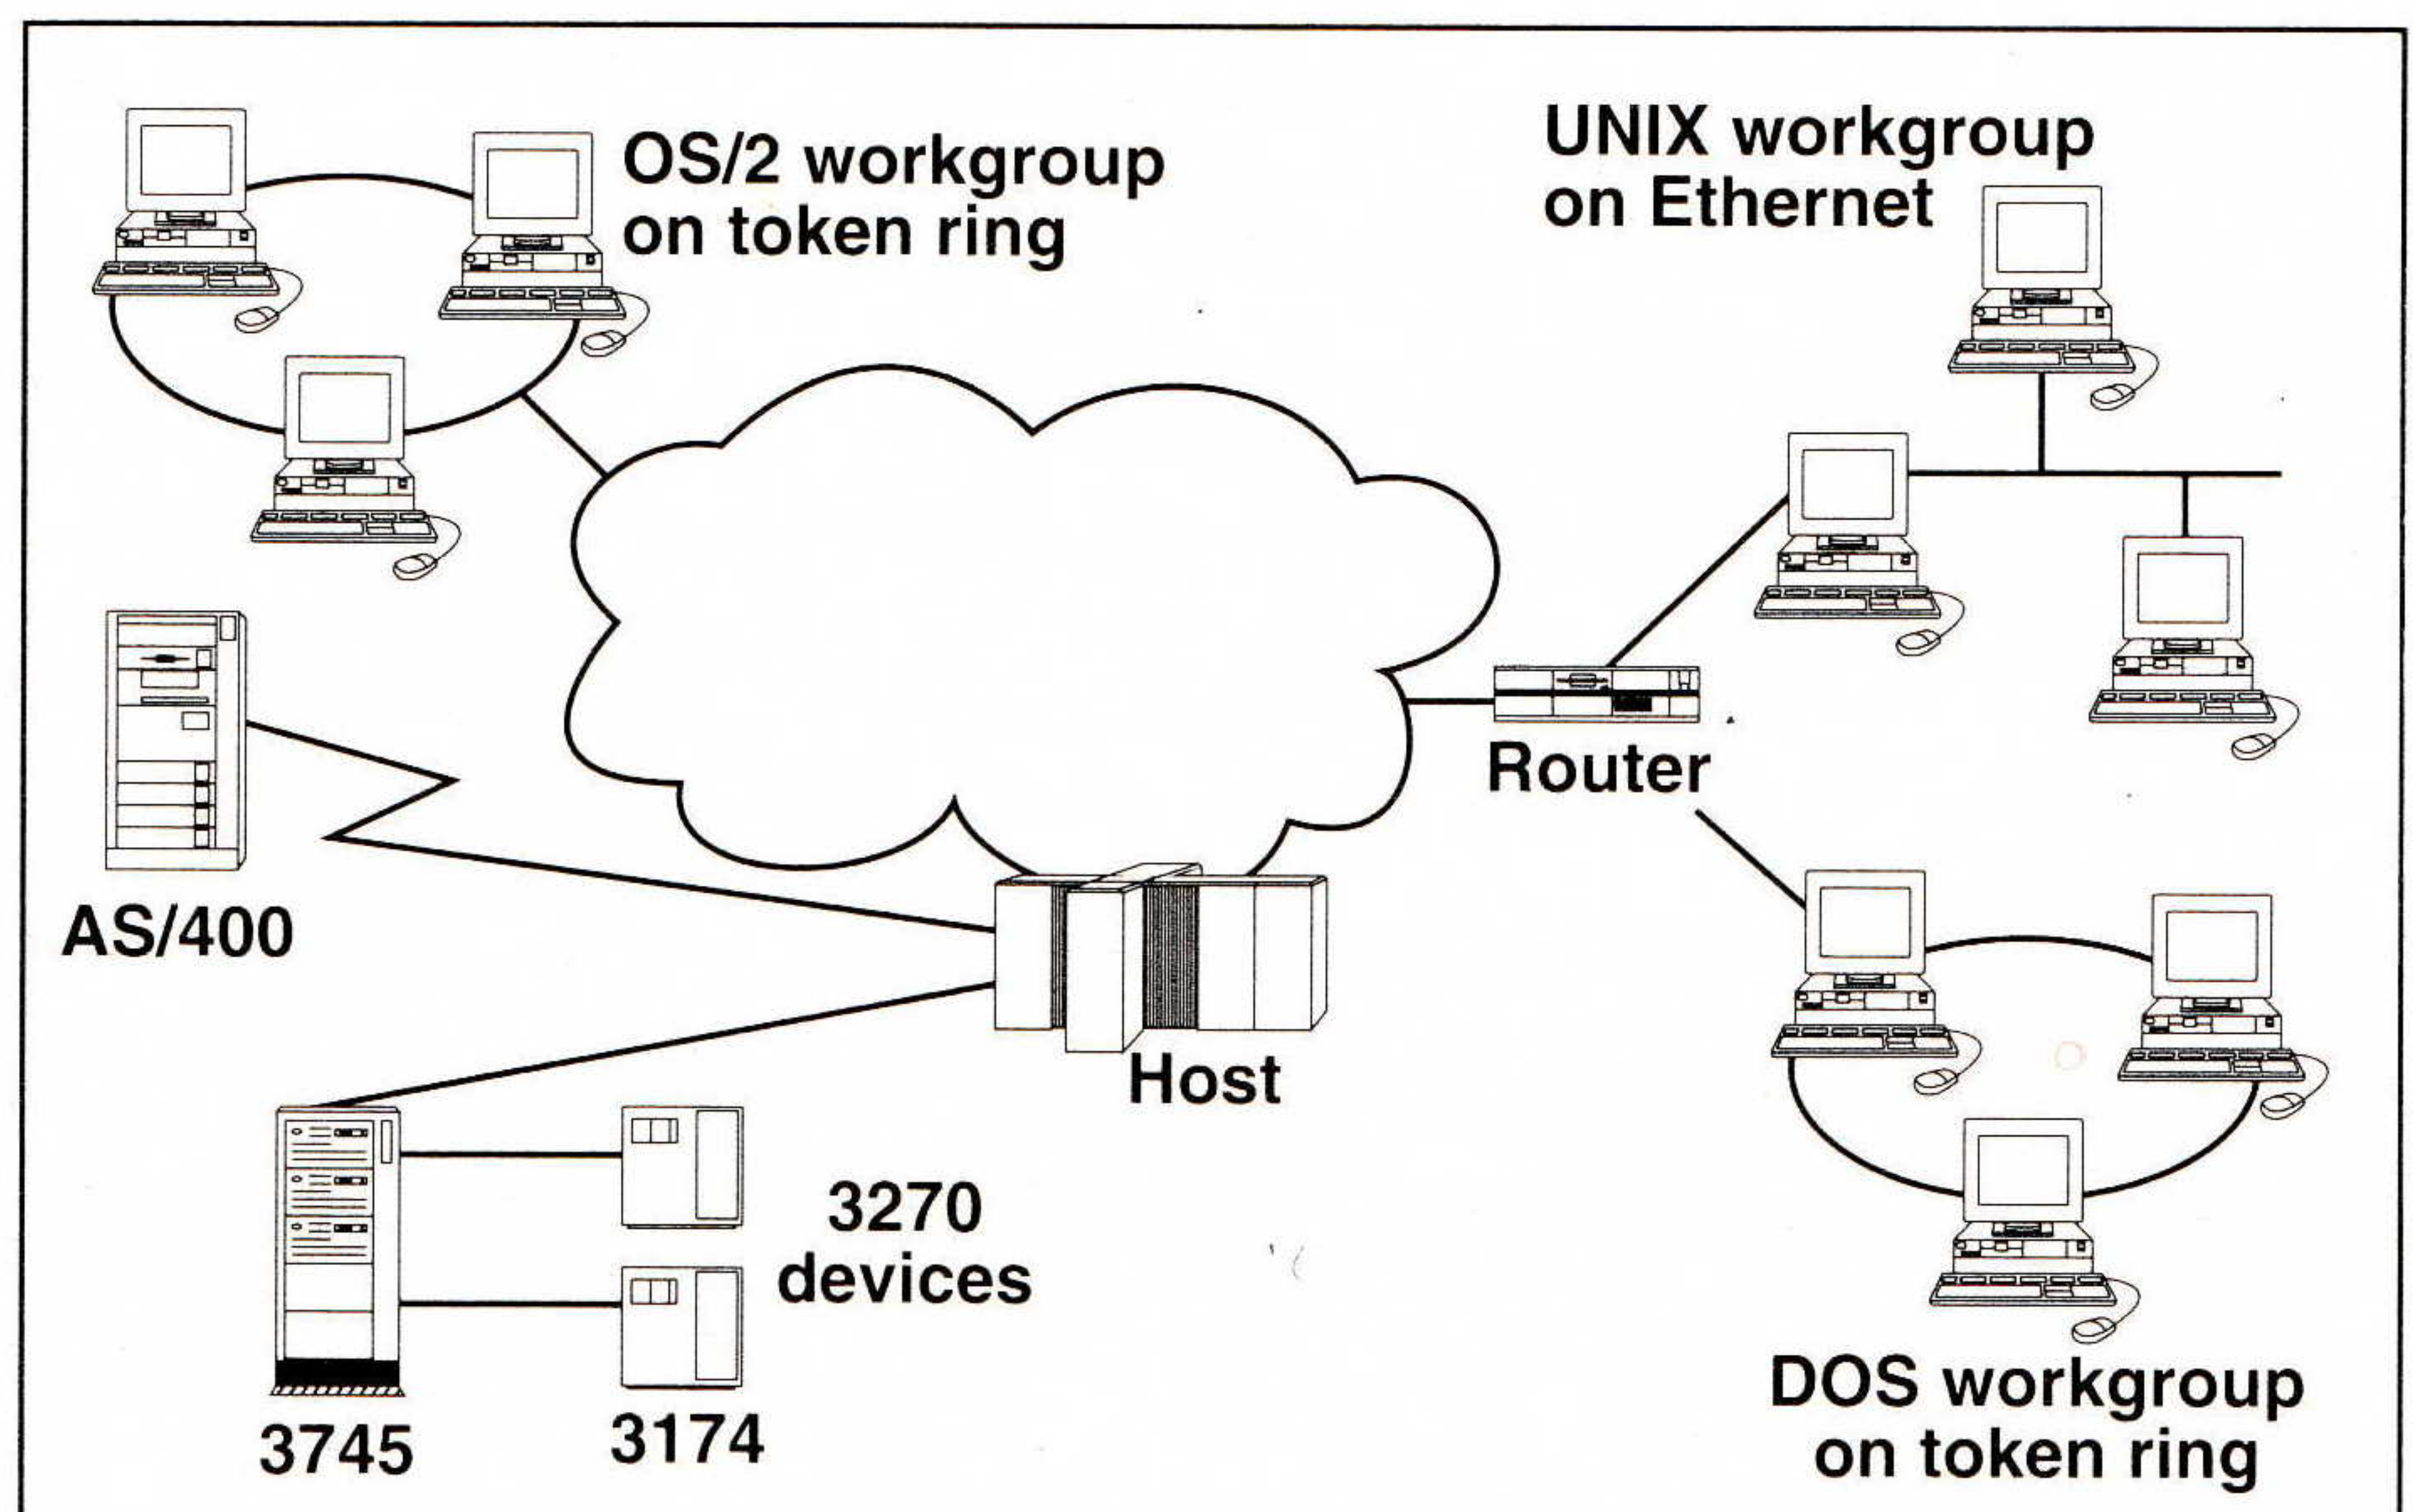


Figure 1: APPN any-to-any connectivity

By using APPC and APPN to integrate your networks, you can overcome the problems with the LAN and WAN approaches:

- *A single network to manage:* With APPC and APPN, you don't have to maintain multiple protocols for different networks. Since you can use a single protocol throughout your network, network integration is easier and less costly.
- *Multiple computers:* APPC and APPN are not limited to a specific type of computer or network. APPC and APPN products are available for every computer in your network.
- *Reduced broadcasts:* APPN nodes never broadcast changes to all machines in the network. Instead, only routing nodes exchange network topology information, and they exchange this information only when changes occur in the backbone of the network. This intelligent approach to routing helps eliminate unnecessary network control traffic.
- *No need for time-outs:* APPC offers guaranteed delivery of data between partners. If something should happen to the data as it travels through the network, APPC automatically notifies both end-points. So, APPC eliminates the need for partners to use time-outs to constantly determine whether data was received.
- *Easy configuration:* APPN automates many configuration tasks that are time-consuming, complicated, and error-prone. With APPN, there's no need to coordinate complicated system definitions across networks. Anytime you move a machine in the network, APPN takes care of the configuration for you.
- *No encapsulation:* Because you can use APPC and APPN throughout your network, encapsulation is not necessary. Therefore, you don't incur the overhead of additional header information, and you can use APPN's *class-of-service* features to move important data through the network faster.
- *Choice of network design:* Unlike the old SNA, which limited you to a hierarchical network, APPC and APPN will support a variety of network designs. You can use APPN to create and interconnect a wide array of network topologies: star, mesh, ring, as well as hierarchical.
- *High performance:* Recent tests show that APPC provides exceptional performance in a LAN environment. Likewise, APPN is widely recognized for its ability to provide both high performance and predictable response time for the applications that need it.

Selecting protocols to support business needs

Choosing which communications protocol to use in your network involves more than a comparison of technical features. It's important to consider how well a particular protocol can help you meet your business needs. Some important considerations include:

- Types of applications and data your users need to access
- Availability of products that support the protocol
- Cost-effectiveness of using the protocol
- Plans for growth of your company and its networks.

Keep in mind that the networking technologies you choose must meet these needs both today and in the years to come.

Using APPC and APPN (*continued*)

Applications for the future

The most important consideration in selecting any networking technology is the needs of your users. What kinds of applications do they need to do their jobs today? What new applications will they need five years from now? You certainly don't want to replace your entire network at some later date to handle needs you can begin to anticipate now. Also, remember that most companies do not migrate their entire staffs to new applications all at once. Because applications are typically phased into an organization over a period of months, you need to be able to continue to use existing applications as you migrate to newer, client-server applications.

These days, many companies have come to depend on relatively new networked applications, like e-mail and file servers, which are well-suited for the LAN environment. Yet, many companies are looking at the next step, *distributed databases*, to address anticipated business needs. Specifically, users need quicker access to large volumes of information, which they need to collate, sort, and filter in real-time. Distributed database applications give these users easy access to data stored on other computers in the network, including more powerful midrange or host machines.

To run distributed database applications, you need a communications protocol that will work with different types of operating systems and computers. Because APPC support is available for a wide range of platforms, it's an ideal choice for distributed databases. In fact, APPC was selected as the communications standard for *Distributed Relational Database Architecture* (DRDA).

Companies are also evaluating more advanced distributed computing applications to take advantage of idle processing power on workstations and midrange machines. By distributing the work between machines, these applications use remote computers to handle processing-intensive tasks like compiling programs and running complex simulations. APPC is available for the various machines that run distributed applications, plus it also provides the exceptional performance required by these applications.

Finally, many companies are already planning for networked multimedia and imaging applications. These applications need both fast physical connections, such as ATM, wireless, Frame Relay, and FDDI, as well as fast networking protocols. APPN is designed to deliver data at a predictable, consistent rate, an important requirement for such applications. In addition, even the fastest, new connection types will continue to need the intelligent routing support that APPN offers. APPN supports all these connection types transparently.

It's not enough to simply provide users with the latest applications. You also need some ways to ensure that only authorized users get access to the data and applications they need. Typically, LAN environments offer only limited security features. By contrast, APPC brings the security features associated with mainframes to all computers in the network. As an added bonus, these security services are consistent across all APPC platforms, so you don't end up with incompatible security systems in your network.

Freedom of choice

Companies used to rely on one or two vendors to supply all of their computing needs. But, these days, no one wants to be locked into using products from a single vendor. Instead, companies want the freedom to choose the best products to meet their business requirements.

Both APPC and APPN give you this freedom of choice. Over 40 vendors, including Novell, Sun, Microsoft, Apple, and IBM, provide platforms or gateways that support APPC. Hundreds of commercially available applications use APPC for communications support, and many more companies are developing custom APPC applications. Similarly, you'll find a number of APPN products available today, and the industry's leading router vendors and platform vendors are adding APPN support to their products.

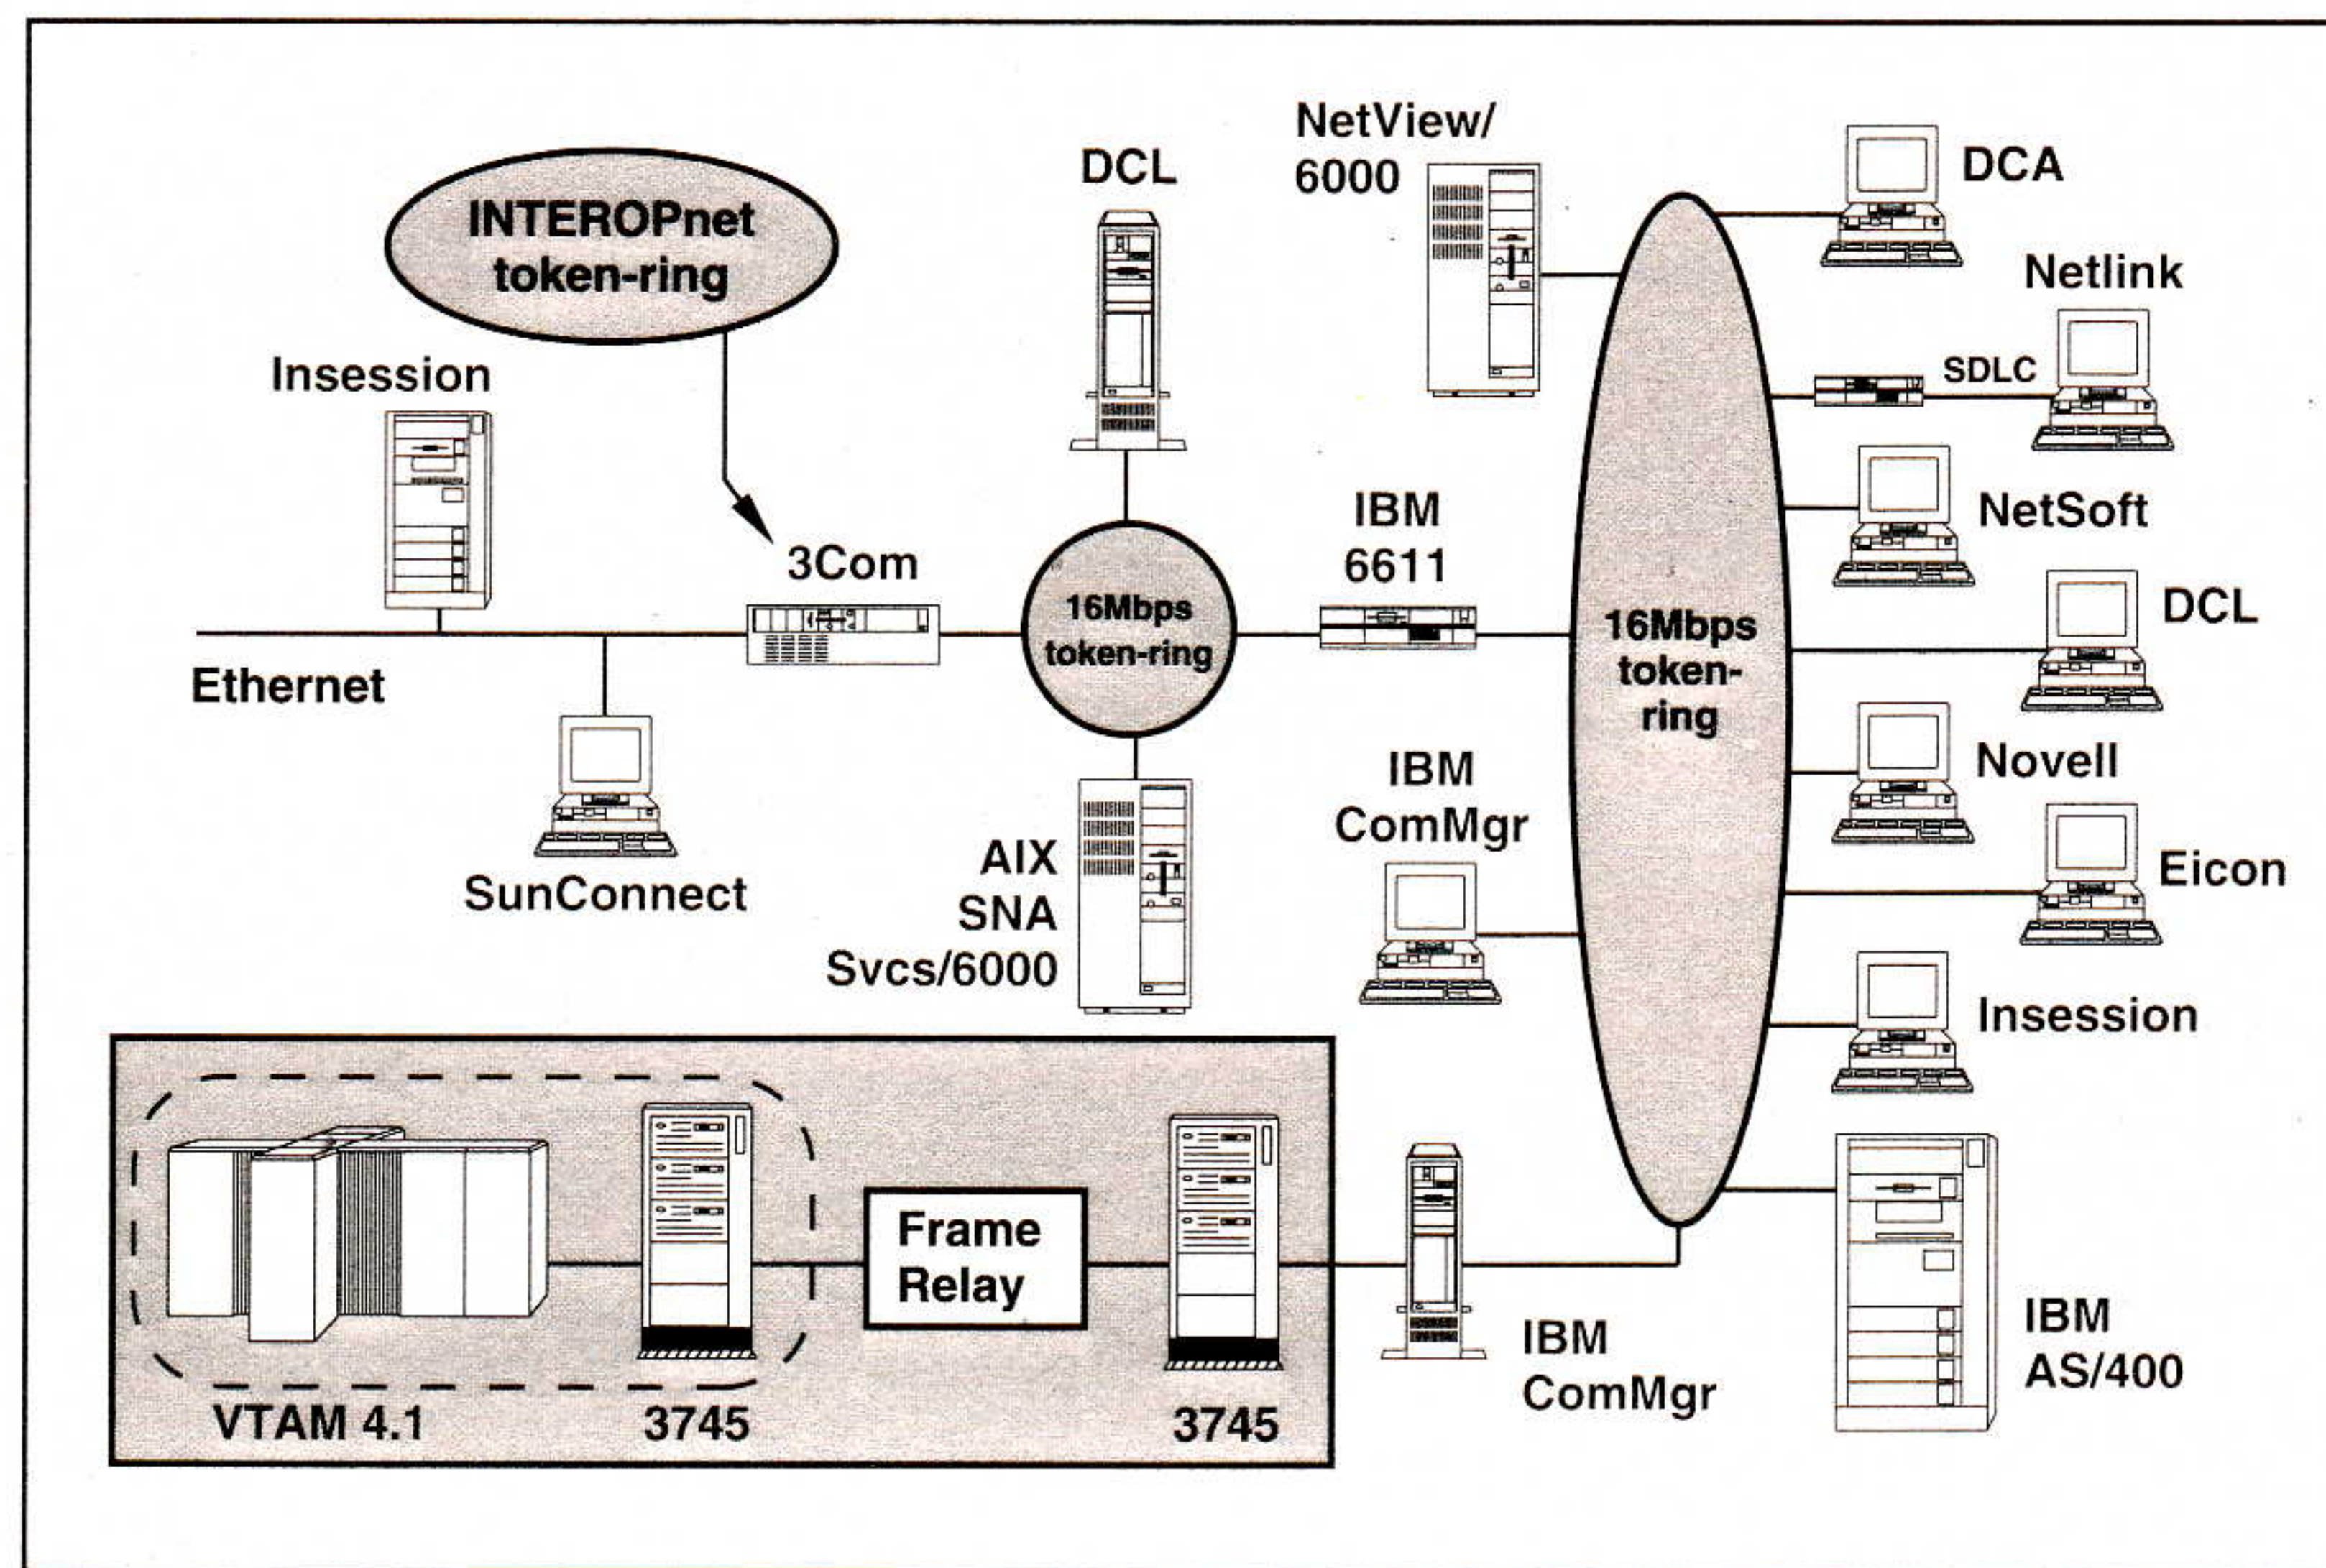


Figure 2: APPC/APPN Solutions Showcase at INTEROP 93 Spring

Cost-effective use of resources

Most companies want to make the most of their existing resources, both equipment and personnel. It's simply not cost-effective to add new lines to the network or devote a full-time administrator to configuration changes. APPN offers several features that help reduce these costs.

As you start integrating LANs and attaching more computers to the network, the traffic flowing through the network increases at a phenomenal rate. Soon, network traffic slows and performance lags. You could add additional lines to the network to solve this problem, but doing so will increase your networking costs. Instead, you can use APPN to help eliminate unnecessary network control traffic, which in turn gives you more bandwidth for moving data through the network.

In addition to the built-in limits on broadcast searches, APPN provides other features that improve network performance. Both caching and central directory services limit searches for other computers in the network. For example, if an APPN network node finds the location of a resource, that information is stored or cached and used the next time the resource is requested. If an APPN network contains a central directory server, all broadcast searches for resources go to the central directory first. Broadcasts are sent to the other network nodes only if the central directory does not know about the resource. These APPN features give you more bandwidth for your applications, without the need to invest in new equipment.

APPN also eliminates the need for additional lines for specific protocols. Often, a user on one LAN needs to access information on another LAN, without going through a hierarchical network. Rather than installing a separate line to support the LAN protocol and give the user direct access, you can use APPN to support peer-to-peer communications between the machines.

continued on next page

Using APPC and APPN (*continued*)

APPN's dynamic configuration features can also free administrators from the complexities of network configuration. For example, if you're installing a new workstation, you simply provide the name of your computer and the address of the intermediate network node that handles your traffic. APPN automatically takes care of the rest of the information needed to route traffic to and from the workstation. Anytime you move a machine, APPN handles configuration changes dynamically. And, with APPN, you don't have to stop the entire network to make these kinds of changes.

As your network grows, the configuration tasks that seem small today may become overwhelming tomorrow. In larger networks, there is a greater likelihood of configuration errors, and even small mistakes in network configuration can lead to costly unscheduled outages. No one wants to manually configure or debug thousands of machines in dozens of remote locations. If you use APPN, you can get error-free configuration for all machines in your network, every time. These advantages give network administrators time to concentrate on other, more important work.

Preparing for the future

Smart businesses invest in technologies that let them solve today's problems and handle network growth. They also plan ahead for future technologies. APPN enables you to migrate today to the new client-server environment, and APPN enhancements will give you the state-of-the-art sophistication needed for tomorrow's networks.

Network growth doesn't happen overnight. An ideal networking technology will let you migrate to new applications and grow your network at your own pace. With APPN, you can migrate your network to the client/server environment as slowly or as quickly as you want. The new version of VTAM, for instance, lets you integrate traditional SNA networks with APPN networks, so that you can migrate portions of your existing networks at your own pace. APPN also enables you to run both existing mainframe applications and new client-server applications at the same time, which make migration much easier. APPN today will also support new developments in fast connection types, as well as the slow connections you may need to use during the migration process.

Investing in a networking technology is a long-term decision. As much as possible, you need to be sure that your choice has the power and performance to support the applications and technologies of the future. New APPN enhancements provide the ideal base for tomorrow's networking technologies. With APPN's *High-Performance Routing* (HPR), you'll get even faster support for your high-speed backbone. HPR will also non-disruptively reroute sessions around failed nodes, so that outages never interfere with the applications that you are running. Another feature, border node, will enable you to connect different APPN networks while reducing the number of topology updates that travel through the entire network.

More information

You can get more information about the benefits of using APPC and APPN to integrate your network from the following:

- *APPC Info Exchange Forum on CompuServe*: This forum provides you with information and answers on APPC and APPN, as well as technical papers and sample programs you can download. To get a free introductory membership, phone CompuServe at 1-800-848-8199 or 1-614-457-0802; be sure to ask for representative 337. To access the forum, type **GO APPC** at the **!** prompt.

- *APPC Market Enablement*: This technical support group publishes *THE APPC CONNECTION*, a bimonthly newsletter that features articles on APPC and APPN, from market trends to practical how-to's. The group also published "Networking with APPN," a free booklet that provides a technical overview of APPN features.

- *APPN Architecture and Product Implementations Tutorial*: This book contains a tutorial on APPN as well as an overview of APPN implementations on IBM platforms. To order this book, fax your order to 1-800-284-4721 and specify order number GG24-3669-00.

References

- [1] Clark, Wayne, "Accommodating SNA Peer-to-Peer Networking in a Multiprotocol Environment," *ConneXions*, Volume 7, No. 3., March 1993.
- [2] Joyce, Steven T. and Walker II, John Q., "Advanced Peer-to-Peer Networking (APPN): An Overview," *ConneXions*, Volume 6, No. 10, October 1992.
- [3] Clark, Wayne, "SNA Internetworking," *ConneXions*, Volume 6, No. 3, March 1992.
- [4] "Networks of Small Systems," IBM Publication No. GG66-0216-00, July 1985.
- [5] "AS/400 Communications: APPN Network User's Guide," IBM publication number SC21-8188.
- [6] "Networking Services/2 Installation and Network Administrator's Guide," IBM publication number SC52-1110.
- [7] "APPN Architecture and Product Implementation," IBM publication number GG24-3669.
- [8] Meindl, R., "Establishing Workgroups in a Multiprotocol Environment" *ConneXions*, Volume 7, No. 10, October 1993.

Learn more:
Tutorial T71
and
Session IT-1.

TIM HUNTLEY is a senior associate programmer with IBM APPC Market Enablement in Research Triangle Park, NC. He is currently working with several vendors on APPC3270, a new utility that enables devices that emulate 3270 terminals to send data encapsulated in an APPC session.

SUSAN SCHULKEN is the managing editor of *THE APPC CONNECTION*, a free bimonthly newsletter on APPC, APPN, and CPI-C. To get a free subscription to the newsletter, send your request (please print!) to: *THE APPC CONNECTION*, Dept. E42/502, P.O. Box 12195, Research Triangle Park, NC 27709. You can also fax requests to 1-919-254-6050.

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Need a *ConneXions* binder? Want to enquire about back issues? (there are now over 80 to choose from; ask for our free 1987-1992 index booklet and the 1993 partial index sheet). We want to hear from you. Send your questions, comments or suggestions to:

ConneXions—The Interoperability Report
480 San Antonio Road
Mountain View, CA 94040-1219
USA

Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)

Fax: +1 415-949-1779

E-mail: connexions@interop.com

A Look at the Host Resources MIB

by Steven L. Waldbusser, Carnegie-Mellon University

[Ed.: This article is reprinted with permission from *The Simple Times*, Volume 2, Number 3, 1993. For more information about *The Simple Times*, see announcement immediately following this article].

Introduction

As the Host Resources MIB nears the IETF standards track, it is constructive to evaluate what it is and how it promises to improve systems management of internet hosts.

The *Simple Network Management Protocol* (SNMP) has been implemented widely on many varieties of platforms, from routers and hubs to workstations and PCs. To date, the standard MIB modules implemented on hosts have been oriented towards the networking aspects of those hosts. The only systems management MIB modules for hosts have been proprietary. Nevertheless, these proprietary modules have been written for many of the host types available such as PCs, Macs, and workstations.

After a fair amount of experience had been gained with these proprietary MIB modules, a desire grew for a standard module which would incorporate the functions common amongst the various vendor-specific modules. Such a MIB would better allow third-party applications to be written that could manage various types of internet hosts. Given that most network installations are multi-vendor and have several different types of host systems, a common MIB module and common applications would allow a network operations center to manage these various types of hosts with a common user interface.

Enough desire grew for this standard host MIB that an IETF working group was formed. That working group has defined a MIB, the *Host Resources MIB*, soon to be evaluated for the standards track.

The Host Resources MIB Module

The MIB is divided into six groups of objects: the system group, the storage group, the devices group, the running software group, the running software performance group, and, the installed software group.

The *system group* has several objects that describe overall system parameters such as the number of users and processes, and where the operating system is loaded from.

The *storage group* provides utilization information on all forms of storage on the system, including RAM, disk drives, and, memory buffers.

The *devices group* provides configuration and fault information about all devices on the system. Some types of devices such as printers and disk drives have more detailed information specified than others. For example, one kind of device is a network interface. MIB-II already defines extensive information on interfaces, so the devices group really doesn't need to duplicate those features.

The *running software group* lists the software running on the system while the *running software performance group* provides performance data about each piece of software running on the system. The *installed software group* lists all the software installed locally on the system.

Expected uses

MIB modules should be written to provide variables with known uses rather than providing all possible information in a specific area without regard to its usefulness. The Host Resources MIB was written with this in mind—in fact, some variables were dropped from early versions of this MIB module because their usefulness could not be proven.

The Host Resources MIB provides many types of functions, primarily in the fault management, configuration and asset management, and, performance management areas.

The Host Resources MIB helps a number of system management jobs that a network manager might face. With this MIB module, a network manager can download an inventory of all equipment on various LANs across an organization, without regard to what types of systems the equipment is attached to. In addition to determining how much memory and disk is installed in each computer, the types and versions of other hardware and software components can be retrieved. Obsolete versions of software or hardware can be flagged and incompatibilities between various hardware and software components can be detected. Disk drives can be monitored to make sure that routine backup procedures are being followed and that the disks are not running out of space.

Textual information on a system, such as that found in file names, is returned in a format that allows any international language or character set. This allows the MIB module to monitor systems that have been internationalized, and (again) shows that SNMP is suited for this task.

A typical fault diagnosis exercise

In a typical environment, the network/systems manager might get a call from a user on a PC or workstation complaining of difficulty running an application. The manager can retrieve information over the network to diagnose the problem. Initially, the manager might check the RAM, disk and buffer usage on the system to see if any allocation failures have been experienced or if any are near their limits. The manager can check to see if too many users are on the system, too many processes are in use, or if the system is otherwise overloading the CPU. For each device, the manager can check the status and if any errors have occurred that might be the cause of the user's problem—for example, printers indicate whether they need paper or any other attention.

After exploring each of these areas of the system, if the problem has not been found, it is likely a software problem. Software systems are becoming more and more complex, and it is especially important to verify the compatibility amongst various pieces of software, such as the operating system, the windowing system, the network file system, and, the applications software. The Host Resources MIB allows the manager to remotely determine what versions of these pieces of software have been installed, and what versions are running. This allows useful applications to be written that know of these incompatibilities and are smart enough to automatically flag such errors. If a piece of software that is running is causing a problem, the manager can use the MIB (with SNMP's security) to kill that process.

Conclusion

The Host Resources MIB provides a solution to many problems that once required vendor-specific MIB modules. Now that this soon-to-be standard MIB module has stabilized, those functions can be implemented in a vendor-neutral manner. In addition to the advantages this provides to application builders, users are provided with a common interface for performing systems management across multiple types of host platforms. This common interface will prove to be useful to solve a variety of problems users have.

*Learn more:
Tutorial T51
and
Session NM-4.*

STEVEN WALDBUSSER is the Manager of Network Development at CMU, where he is responsible for providing network management systems to manage a heterogeneous campus internetwork of 6000 hosts and 150 networks. He is a member of the IETF's Network Management Directorate and the author of several RFCs. E-mail: sw01+@andrew.cmu.edu.

Information on *The Simple Times*

Contents

The Simple Times (ISSN 1060-6068) is an openly-available publication devoted to the promotion of the *Simple Network Management Protocol*. In each issue, *The Simple Times* presents: a refereed technical article, an industry comment, and several featured columns:

- Applications and Directions
- Security and Protocols
- Working Group Synopses
- Ask Dr. SNMP
- Standards

In addition, some issues include brief announcements, summaries of recent publications, and an activities calendar. Past technical articles have included:

- A New View on Bulk Retrieval with SNMP
- Sets are Fun: Introducing the SMDS Subscription MIB Module
- An Implementation of SNMP Security
- Accomplishing Performance Management with SNMP
- An Introduction to SNMP MIB Compilers

Policy

The Simple Times is openly-available. You are free to copy, distribute, or cite its contents. However, any use must credit both the contributor and *The Simple Times*. Further, this publication is distributed on an "as is" basis, without warranty. Neither the publisher nor any contributor shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused, directly or indirectly, by the information contained in *The Simple Times*.

Subscription information

The Simple Times is available via electronic mail in three editions: *PostScript*, MIME (the multi-media 822 mail format), and richtext (a simple page description language). For more information, send a message to: st-subscriptions@simple-times.org with the word "help" in the Subject: field. In addition, *The Simple Times* has numerous hard-copy distribution outlets. Contact your favorite SNMP vendor and see if they carry it. If not, contact the publisher and ask for a list:

The Simple Times
c/o Dover Beach Consulting, Inc.
420 Whisman Court
Mountain View, CA 94043-2186
Tel: +1 415-968-1052 • Fax: +1 415-968-2510
E-mail: st-editorial@simple-times.org

Submission information

The Simple Times solicits high-quality articles of technology and comment. Technical articles are refereed to ensure that the content is marketing-free. By definition, commentaries reflect opinion and, as such, are reviewed only to the extent required to ensure commonly-accepted publication norms. *The Simple Times* also solicits announcements of products and services, publications, and events. These contributions are reviewed only to the extent required to ensure commonly-accepted publication norms.

Submissions are accepted only in electronic form. A submission consists of ASCII text. (Technical articles are also allowed to reference encapsulated *PostScript* figures.) Submissions may be sent to the contact address above, either via electronic mail or via magnetic media (using either 8-mm *tar* tape, 1/4-in tar cartridge-tape, or 3-1/2-in MS-DOS floppy-diskette).

Each submission must include the author's full name, title, affiliation, postal and electronic mail addresses, telephone, and fax numbers. Note that by initiating this process, the submitting party agrees to place the contribution into the public domain.

Call for Papers

The *Journal of High Speed Networks* (JHSN) announces a forthcoming issue in early 1994 on *Quality of Service* (QOS) in Integrated Services Networks. JHSN publishes high quality papers on a number of topics ranging from design to practical experience with operational high speed networks.

Topics

An important benefit of high speed networking technology is its ability to support a wide range of applications. However, the extent to which the resources of a store-and-forward network are shared may compromise Quality of Service in a number of different ways. While there has been considerable research in the area, there remain many fundamental issues that need to be better understood. We solicit high quality papers that address these issues either through analytical or experimental techniques. Representative areas of interest include:

- Appropriate metrics for QOS;
- Application requirements in terms of network performance;
- Experimental results on human perception of network performance degradation through loss and delay;
- Network Services to better support distributed multiuser applications.
- Evaluating current proposals for QOS against specific application requirements;
- Performance models for analyzing QOS;
- The nature of QOS guarantees offered by the network i.e., firm, best-effort or something in between;
- Preserving QOS across heterogeneous network domains;
- Network Control strategies to facilitate QOS.
- The tradeoff between supporting diverse applications and simplifying network control mechanisms;
- The tradeoff between supporting diverse applications and enabling efficient use of network resources

Submissions

Send 5 copies of the submission to:

Dr. Abhay K. Parekh, Guest Editor
IBM
T. J. Watson Research Center
PO Box 704
Yorktown Heights, NY 10598
E-mail: parekh@watson.ibm.com
Voice: 914-784-7888

The deadline for submission is November 30, 1993.

More information

For more information on the *Journal of High Speed Networks*, please contact the Editor-in-Chief:

Professor Deepinder Sidhu
Maryland Center for Telecommunications Research
University of Maryland – BC
Baltimore, MD 21228-5398
E-mail: sidhu@umbc.edu
Voice: 410-455-3028
Fax: 410-455-3969



A Further Step in Interop Company's Worldwide Growth

About Interop Company

Interop Announces NetWorld+Interop 94 Tokyo

Interop Company recently announced *NetWorld®+Interop® 94 Tokyo*, a combination conference and exhibition focused on the technical and educational aspects of integrating all levels of communications and computing. The new event will be held July 25–29, 1994 at the Makuhari Convention Center in Tokyo, Japan. The goal of the new conference and exhibition is to provide computer and communications professionals in business and government with an educational forum for learning about the products, technologies and interoperability issues critical to success in today's global economy.

NetWorld+Interop 94 Tokyo is Interop Company's first major event in Japan and the Asia Pacific region. Like all Interop events worldwide, the conference portion of this event will be managed by a program committee of local industry experts with in-depth knowledge of networking standards, protocols and other relevant technology issues.

This announcement of a Tokyo venue is another step in plans first unveiled by Interop Company and Novell, Inc. on December 18, 1992. At that time, the two companies announced the creation of an event called NetWorld+Interop 94, a combined conference and exhibition for addressing the industry's growing need worldwide to understand interoperability issues at all levels of networking and computing—from the desktop to the data center.

"Trade shows and academic conferences are usually held separately and target different audiences," said Dan Lynch, chairman of Interop Company. "That situation is changing, however, as technology becomes more important and business and government leaders around the world realize that interoperability is integral to their computing and networking purchase decisions. They seek a forum where they can study the complex issues of internetworking, learn about the available solutions and see how different vendors' products work together. In the past year, we have seen a growing number of international visitors at Interop Company events in the U.S., asking when we plan to bring the 'Interop experience' closer to their home territories. Tokyo is an important step in satisfying that request."

Interop Company is a conference and exhibition company specializing in computers and communications. Interop is dedicated to creating new standards in expositions worldwide that are rich in content and concentrate on technically excellent demonstrations. Interop events are designed to bring computer and communications professionals the latest ideas from researchers, analysts and vendors through conferences, seminars, publications and educational services. Founded in 1985, Interop Company is the sponsor of conferences and expositions worldwide, including NetWorld+Interop 94 and INTEROP, the premier forums for addressing the interoperability challenges and solutions found in the real world of enterprise computing, from the desktop to the data center.

Interop Company is headquartered at 480 San Antonio Road, Mountain View, California 94040–1219. Interop Europe is located at Immeuble Omega, 10 rue Thierry le Luron, 92593 Levallois-Perret Cedex, Paris, France. Interop Japan is located at ABS Bldg. 202, 2–4–16, Kudan-Mianmi, Chiyoda-ku, Tokyo 102, Japan.

Interop is a registered trademark of Interop Company; INTEROPnet is a trademark of Interop Company; NetWorld is a registered trademark of Novell, Inc. Bruno Blenheim Inc. is not affiliated in any way with NetWorld+Interop 94.

Vive La NetWorld+Interop 94 Paris Conference!

To help address the growing need for global computing in France, Interop Company announced recently its fifth NetWorld®+Interop® conference, *NetWorld+Interop 94 Paris*, a combination conference and exhibition focused on the technical and educational aspects of integrating all levels of communications and computing. This event will be held October 24–28, 1994 at CNIT, La Défense in Paris, France. The four other venues for NetWorld+Interop 94 include: Las Vegas, Nevada, May 2–6; Berlin, Germany, June 6–10; Tokyo, Japan, July 25–29; and Atlanta, Georgia, September 12–16, 1994.

Worldwide demand for interoperability

NetWorld+Interop 94 Paris will be the next French computing revolution after the upcoming INTEROP Europe 93 Paris conference to be held October 25–29 at CNIT, La Défense, Paris. More than 20,000 attendees and 200 exhibitors are expected at the INTEROP Europe 93 Paris conference. The follow-on goal of NetWorld+Interop 94 Paris is to provide French computer and communications professionals in business and government with an interactive educational forum for learning about the products, technologies and interoperability issues critical to success in today's global economy.

"In France, as companies address pan-European and global communication requirements, the need for application networking and protocol interoperability is growing dramatically," said Daniel C. Lynch, chairman and founder of Interop. "This worldwide interoperability requirement is the successful driving force behind NetWorld+Interop 94 in Paris. The projected attendance numbers at this October's INTEROP Europe 93 Paris clearly illustrate the international market demand for a forum to study the complex issues of internetworking, learn about the available solutions, and see how different vendors' products work together."

Unique requirements

The European market has unique computing requirements. These include reconciling international standards, integrating a large amount of foreign technology, and agreeing upon a common language for a multicultural business community. Like all the NetWorld+ Interop 94 events planned in Las Vegas, Nevada; Berlin, Germany; Tokyo, Japan; and Atlanta, Georgia; the Paris conference and tutorial program will be managed by a program committee of Pan-European and global experts with in-depth knowledge of networking standards and technology issues relevant to worldwide audiences. In addition to the exhibition and general conference, the five-day Paris event offers delegates a choice of special conferences, in-depth tutorials led by European and American academic experts, the INTEROPnet™ network and Solutions Showcase™ Demonstrations highlighting the practical applications of new technologies.

Mark your calendar!

The five locations and dates for NetWorld+Interop 94 are:

<i>NetWorld+Interop 94, Las Vegas, Nevada:</i>	<i>May 2–6, 1994</i>
<i>NetWorld+Interop 94, Berlin, Germany:</i>	<i>June 6–10, 1994</i>
<i>NetWorld+Interop 94, Tokyo, Japan:</i>	<i>July 25–29, 1994</i>
<i>NetWorld+Interop 94, Atlanta, Georgia:</i>	<i>September 12–16, 1994</i>
<i>NetWorld+Interop 94, Paris, France:</i>	<i>October 24–28, 1994</i>



Paris

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp.Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

Back issues available upon request \$15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779
connexions@interop.com

CONNEXIONS